Glass Ceilings or Sticky Floors? Statistical Discrimination in a Dynamic Model of Hiring and Promotion

David Bjerk*

david.bjerk@cmc.edu

July 12, 2007

Abstract

I show that when two groups differ in (i) their average skill level, (ii) the precision with which they can signal their skill prior to entering the labor market, and/or (iii) the frequency with which they have the opportunity to signal their skill prior to entering the labour market, then even if firms become increasingly informed regarding each worker's skill over time, equally skilled workers from different groups will have different likelihoods of making it to top jobs in the economy, even though there is no discrimination when it comes to promotion to these top jobs.

JEL classification codes: J31, J71, J78.

Attention Printer: pagehead title "Glass Ceilings or Sticky Floors?"

*Thanks to Tom Crossley, Tilman Klupp, and seminar participants at Cambridge University and Emory University. Much of this work was completed while I was on faculty at McMaster University and an NICHD Research Fellow in Population Studies at the RAND Corporation.

In 1995, 95 percent of senior-level managers in the top Fortune 1000 industrial and 500 service companies in the U.S. were male, 97 percent of which were also white (Federal Glass Ceiling Commission, 1995). In the U.K. females make up just 8.3 percent of the senior judiciary, 9.7 percent of top business leaders, and 9.1 percent of newspaper editors (EOC, 2005). At the end of the 2004 American football season, only 5 of the 32 NFL head coaches were black, and only 2 of the 117 Division I College football teams had black head coaches (Rosenburg, 2004). In 1999, there were only 68 female full professors at the 120 Ph.D. granting economics departments in the U.S. (CSWEP, 1999). This severe underrepresentation of females and minorities at the highest levels of occupational achievement throughout a wide range of occupations has often been referred to as the "Glass Ceiling".[1]

It is sometimes argued that the reason why females and minorities are scarcely represented at the top levels of most professions is not because of any explicit discrimination or barriers to promotion to these top jobs, but rather simply because there are too few female and minority applicants with sufficient qualifications. Given the enormous amount of human capital investment required to be qualified for these jobs, including advanced schooling, long hours, and great effort, it is not unreasonable to think that females (possibly due to greater time constraints) and minorities (possibly due to greater monetary and credit constraints) are on average less able or willing to make such investments than white males.[2] Moreover, given the legal and public relations costs associated with being revealed to be an actively discriminating employer, the fact that qualifications and past accomplishments for individuals applying for top positions in the economy are generally quite easily verifiable, and the fact that

[1] For further empirical work on the glass ceiling phenomenon, see Winter-Ebmer and Zweimuller (1997), McDowell, Singell, and Ziliak (1999), and Albrecht, Bjorkman, and Vroman (2003).

[2] See Becker and Tomes (1979), Coate and Loury (1993), and Becker (1979) for further theoretical discussion on how human captial inequality can arise between races and genders.

the hiring process with respect to top occupational levels are often subject to substantial scrutiny by individuals both within and outside of a particular institution, it is difficult to believe that direct discrimination in the hiring and promotion process for the top jobs is the dominant reason for the tremendous underrepresentation of females and minorities in these jobs.

However, in this paper I show that even if there is no discrimination with respect to promotion to the top jobs in the economy, the reasons for the glass ceiling phenomenon may be more complicated than there just being too few females or minorities who can do the job. Indeed, I show that underrepresentation of one group relative to another (e.g. females vs. males, blacks vs. whites) in the top jobs in an occupation can arise not only if one group has higher average skill than the other, but also if firms are less confident in the precision of the pre-market and early market skill signals they observe from members of one group relative to the other, and if members of one group generally have fewer opportunities to signal their skill prior to labour market entry (or early in their career) than members of the other. Moreover, if *any* of the above inter-group differences hold, then members of one group will be less likely to make it to the top of their profession than *equally capable* members of the other group, even when there is no discrimination with respect to promotion to the top jobs. In other words, while there may be no direct discrimination with respect to hiring or promotion into the top jobs in the economy, there may still be inequality of opportunity with respect to getting there.

I formalize the above assertions below by developing a dynamic model of the labour market where workers can differ in their skill type, and workers can be hired into or promoted to three job levels in the economy—low level jobs, career-track jobs, and director level jobs—over the course of their career. While at a given job, workers are assumed to be assigned a sequence of tasks.

Successful task completion generates increasing revenue for the firm as the job level rises. However, task failures become increasingly costly to the firm as the job level rises. I assume that all workers, regardless of skill type, will successfully complete all tasks they undertake at low level jobs. However, as the job level rises, less skilled workers become increasingly likely to fail a given task they are assigned at that job.

In this environment, under perfect information, profit maximizing firms would optimally choose to promote only the more highly skilled workers to career-track jobs and especially to director level jobs. However, it is assumed that firms cannot directly observe worker skill. Rather, firms can only observe each worker's track record for successfully completed tasks at each job they have worked at, as well as the skill signals each worker emitted prior to starting his or her career or emitted during the time he or she worked at the low level job.

Given this structure, optimal firm behaviour is to essentially set "qualification" thresholds required for hiring/promotion to career-track and director level jobs respectively. In particular, in order to be hired/promoted to a career-track job, workers have to emit some minimum number of positive skill signals prior to joining the labour market or while they worked at low level jobs. Similarly, in order to be hired/promoted to a director level job, workers have to successfully complete some minimum number of tasks during their time at a mid-level career-track job.

As foreshadowed above, a worker's group affiliation (i.e. gender, race, nationality, etc.) is allowed to enter the model in three distinct ways. First, I allow the fraction of a group that is highly skilled (i.e. sufficiently skilled so as to be a profitable worker at the director level jobs) to differ across groups. Second, I allow the precision of the skill signals emitted prior to or early in workers' careers to differ across groups. And third, I allow the average frequency with which individuals have the opportunity to signal their skill prior to or early in their careers to differ across groups.

I show that either of the first two of these possible differences across groups will cause profit maximizing firms to discriminate by requiring more positive skill signals in order to hire/promote members of one group to the career-level job than members of the other group. However, once at the career-level job, there will be no discrimination with respect to *wage rates* or *promotion to the director level job* (the intuition for this result is discussed in more detail in the literature review below). This greater number of successful skill signals required for hiring/promotion to the career-track job will mean that, on average, it will take individuals from one group longer to be hired/promoted to the career-track job than *similarly* skilled individuals from the other group. Alternatively, the third possible difference between groups (i.e. a difference in the average frequency with which individuals have the opportunity to signal their skill prior to or early in their careers) will not lead to any explicit discrimination in the model. However, such a difference in frequency of opportunity for signalling skill will still mean that it will on average take individuals from one group longer to be hired/promoted to the career-track job than similarly skilled individuals from the other group, as it will take them longer to emit any given number of positive skill signals that is required for hire/promotion to a career-track job.

This longer average time to hiring/promotion to the career-track jobs is of considerable interest. In particular, in the model there is no explicit discrimination with respect to promotion to the director level jobs, indeed once hired to the career-track job everyone regardless of group will have to successfully complete the same number of tasks in order to be hired/promoted to the director level jobs. However, the delays in obtaining a promotion to the career-track for members of one group will mean that a smaller fraction of individuals from this group will have the time to develop this required success history while at a career-track job during the remainder of their careers than *equally skilled* members of the other group. This implies that even when firms do not inherently "prefer" to hire members of one group over members of another, there can still be inequality of opportunity between groups of reaching the highest job levels in an occupation. In this way, this model suggests that underrepresentation of females and minorities among those working in at the highest level may not be so much of a "glass ceiling" but rather a "sticky floor."

1 Related Literature

This model is certainly very closely related to the "classical" models of statistical discrimination in wages (examples include Phelps, 1972; Aiger and Cain, 1977; Lundberg and Startz, 1983; Coate and Loury, 1993; Moro and Norman, 2004). Generally, these models of statistical discrimination are static games, where firms cannot directly observe a worker's skill, but are able to observe some piece of information about a worker that can be used to estimate a given worker's skill. One reasonable question to ask regarding this literature is how the results change when the environment is modelled in an arguably more realistic dynamic setting, where employers progressively learn more about each worker's skill over the course of his or her career.

A primary way in which the model presented below extends the statistical discrimination literature is to incorporate such a dynamic element. In this way, this paper is similar to Fryer (forthcoming), who also considers statistical discrimination in the context of a dynamic rather than static setting. Fryer provides an interesting argument showing that when there is a hierarchy of jobs within a firm, a firm discriminating against one group with respect to hiring into a lower level job *may* switch to discriminating against the other group when it comes to promotion to a higher level job. The intuition being that if the firm requires a higher standard for one group at the initial hiring stage, it may then be able to be more lenient at the promotion stage for this group since it is more confident in the ability of workers' from this group once hired.[3]

Like Fryer's model, the model developed below also shows that in a dynamic setting, discrimination in promotion may arise but in a non-uniform way across the job ladder. Specifically, as discussed in the introduction, profit maximizing firms may choose to discriminate with respect to promotion/hiring into the mid-level career-track jobs, but not discriminate with respect to promotion/hiring into the highest level director jobs. The intuition for this result is as follows. In the model below, a firm discriminates by requiring workers from one group to have more positive skill signals in order to be promoted/hired to the career-track job than workers from the other group. However, by promoting/hiring an individual to the career-track job, a firm must implicitly believe a worker is of a sufficiently high type such that it will be profitable to hire/promote him or her, regardless of his or her group. Therefore, once hired to the career-track job, all that matters for promotion to the director job is the worker's performance at tasks in the career-track job. Therefore, in contrast to Fryer's switching equilibrium, the group discriminated against in lower level jobs in this model will be unambiguously worse off than the other group, since firms never find it optimal to reverse discriminate with respect to promotion in higher level jobs.

There are also several other differences between the model proposed here and that proposed by Fryer. First, his model is explicitly only a two-period model in which employers only observe information about each worker twice. By contrast, the model presented below allows for an arbitrarily large number

[3]It should be noted that this specific result is not general. Indeed, as Fryer points out, several conditions must be imposed on the model's parameters for such an equilibrium to exist, even in his proposed environment.

of periods, as well as an arbitrarily large number of times employers can observe information about each worker's skill.

Another key difference between this model and that developed by Fryer is that in the model presented here, a worker's type is exogenously given prior to the worker joining the labour force, where this type is what determines whether a worker is qualified for the upper level jobs. Firms then use information they observe over the whole course of the worker's career to attempt to infer this type. By contrast, in Fryer's model, workers become qualified for the next level job only by making costly investments each period. Moreover, after hiring a worker, Fryer assumes a firm can perfectly observe whether that worker was actually qualified for the initial job level. This means that in Fryer's model, when deciding whether to promote workers to the high level jobs, firms do not care about earlier signals the worker emitted. Rather, they only care about the skill signal the worker emits subsequent to being hired for the low level position, since that is the only one that contains the relevant information regarding whether the worker made the appropriate investment to be qualified for the high level job. While the endogeneity of skills/qualification is an important feature of Fryer's model that is not contained in the model presented here, employers in Fryer's model do not necessarily become increasingly more informed about a worker's unobserved skill type over time as they do in the model below.

Finally, in both Fryer's model and the model presented below, discrimination occurs with respect to hiring/promotion rather than wages. In Fryer's model, however, this is a direct result from the implicit assumption that wages are exogenously determined. If firms could wage discriminate in his model, they may then choose not to discriminate in promotion. Alternatively, in the model below, wages are determined endogenously, but firms *find it optimal* to discriminate in *hiring* rather than in *wages*.

The model presented here also builds on the classical statistical discrimination literature in that it shows how three distinct possible differences between groups can result in inequality across groups, with inequality becoming even more pronounced if more than one of these differences come into play (which, as I argue below, is certainly plausible). While differences in average skill level across groups and differences in skill signalling precision across groups have been much discussed in the statistical discrimination literature, different frequencies of opportunity for signalling skill across groups has seen very little formal discussion. Moreover, while previous work has generally looked at the labour market effects of differences between groups in only one dimension in isolation, this model shows explicitly how inequality between groups can be further exacerbated when groups differ in more than one dimension.

This paper is also bears some relation to Lazear and Rosen's (1990) gender specific model of statistical discrimination.[4] They assume that each worker's productivity is initially unobserved by firms at the time of hire, but after hiring, firms observe each worker's skill and then decide whether to promote each worker to a higher earnings job track. The key assumptions of their model are that firms lose money if they promote workers who subsequently quit, and, on average, females place a higher value on doing homecare. This causes females to be more likely to quit, all else equal, causing firms to be less likely to promote females in equilibrium.

[4] Another paper of related interest is Athey, Avery, and Zemsky (2000), whose work is also applicable to understanding glass ceiling type phenomena. However, their focus is on how mentoring affects a worker's human capital, where a worker gets more productivity benefits from mentoring, the more mentors there are from his or her group in the upper ranks of the company. In the model developed here, a worker's productive skill (i.e. human capital) is assumed to be fixed, so there are no productivity gains that arise when a worker works for a firm that has more members of his own group in upper management. However, in this model, diversity in management may still be important with respect to mentoring, as management mentors from the underrepresented group may allow low-level workers from the underrepresented group more opportunities to signal their skill than would arise if fewer members of upper management came from the underrepresented group. This issue will be discussed in more detail in Section 6.

While this model has a number of interesting features and results (including some which are somewhat similar to those arising in this paper), the motivation for their assumption that females place a higher value on homecare is not perfectly clear.[5] Moreover, the motivation for their assumptions mean the Lazear and Rosen model is only relevant for thinking about glass ceiling type phenomena with respect to gender. By contrast, the glass ceiling model presented below can be plausibly applied to a wide variety of group distinctions beyond gender including race, nationality, and others.

2 Model

Assume workers are of two skill types, type-h and type-ℓ, from two directly observable groups a and b, and each work/live for a total of T periods. Further assume that there are many firms in the economy, each with three job levels—low level jobs (job $j = 0$), career-track jobs (job $j = 1$), and director level jobs (job $j = 2$). At each job, workers perform one task each period they work. A successful completion of a task at job j generates revenue to the firm equal to $R_j > 0$, while failure at a task at job j generates a loss to a firm of $-L_j < 0$. Assume that $L_j - L_{j-1} > R_j - R_{j-1} > 0$ for $j = 1, 2$. This expression means that the size of the loss associated with a failure relative to the gain associated with a success is assumed to be larger at the higher level jobs. Moreover, since both of the difference terms in the previous expression are assumed to be positive, the above expression also says that failures at higher level jobs lead to greater losses than do failures at lower level jobs, and that successes at higher level jobs generate higher revenue than do successes at lower level jobs.

[5] Specifically, it is not precisely spelled out how this homecare is implicitly valued. If the value of time doing homecare is assumed to be the cost of purchasing such services from the market, then it is unclear why women would face higher costs for purchasing such services then men, when a household most often contains a man and a woman. This seems to be the implicit assumption since workers are assumed to work if and only if their wage exceeds this "value" of their homecare services.

Worker productivity at these jobs is related to worker type in the following manner. A type-h worker succeeds at a job j task with probability one. Alternatively, a type-ℓ worker succeeds at a job j task with probability $1 - \pi_j$, but fails with probability π_j, where $\pi_2 \geq \pi_1 > \pi_0 = 0$.[6] In other words, there is some positive probability of failure for type-ℓ workers at job 1 and 2 tasks, but not at job 0 tasks. Let the fraction of workers from group g (for $g = a, b$) who are of type-h be denoted α_g.

Assume that a worker's success and failure history for each task he or she has performed at each job he or she has worked is publicly observable for the whole of the worker's career. However, the key assumption of the model will be that firms cannot directly observe each worker's actual skill type. Rather, in addition to the successes and failures at previous job tasks, employers can observe whether a worker emitted a positive or negative skill signal on each skill signalling opportunity the worker has had. Assume that prior to joining the labour market or while working at the low level job, the probability that a worker from group g (for $g = a, b$) has a skill signalling opportunity in a given period equals $\phi_g \in (0, 1]$. Note that since everyone succeeds at job 0 tasks, observing the result of a skill signalling opportunity is the only way employers can learn about the skill type for individuals who have yet to join the labour force or for individuals who have only worked at the low level job. Once reaching job level 1 or 2, however, a worker's success or failure record at the tasks he or she is assigned is informative, but it will be assumed that a worker no longer has separate opportunities to signal his or her skill.

Skill signals are informative in that the probability that a type-h worker emits a positive skill signal given a signalling opportunity is equal to one, while

[6] Note that all of the results in this model will generally hold even if type-h workers also failed at tasks with some positive probability, as long as the probability that a type-h worker fails at a task at any given job is less than the corresponding failure probability for a type-ℓ worker.

the probability that a type-ℓ worker emits a positive skill signal given the opportunity is equal to $1 - \lambda_g < 1$. Clearly, the greater λ_g, the more precise the skill signal. Finally, say $\pi_1 > \lambda_g$ (for $g = a, b$), meaning whether or not a worker succeeds or fails at a job 1 (or 2) task is a more precise predictor of his or her type than a positive skill signal.

Groups are allowed to differ in three distinct ways. First, α_a may be greater than α_b, meaning members of group a might be more skilled on average than members of group b.[7] Second, λ_a may be greater than λ_b, meaning skill signals may be more precise for members of group a than members of group b. This is a common assumption made in models of statistical discrimination (Aiger and Cain, 1977; Lundberg and Startz, 1983; Lang, 1986; Bradford and Welch, 1996), and is meant to capture the notion that employers from one group may be less accurately able to assess skills for members from another group (where it is assumed that those making hiring decisions are generally from group a). In the context of this model, if we think of group a as males and group b as females, there are several reasons why it might be true that $\lambda_a > \lambda_b$. For example, psychologists often point to gender differences in communication style (Tannen, 2001), which may cause male superiors to less precisely assess ideas coming from female subordinates than male subordinates. Similarly, females may often participate in different activities from males while growing up (e.g. dance vs. athletics), and male employers may have more difficulty assessing how the skills learned in these less familiar activities may translate to the work world.

Third and finally, ϕ_a may exceed ϕ_b, meaning members of group a may have more frequent opportunities to signal their skill than members of group b.

[7] The results in Coate and Loury (1993), Moro and Norman (2003), and Fryer (forthcoming) are all driven by overall skill differences between groups. However, unlike these papers where such skill differences arise endogenously from similar underlying population, any skill differences between groups in this paper are taken as exogenous. More precisely, the underlying process that determins skill is simply left unmodelled, meaning I only model the labor market consequences of differences in skill distributions that exist between groups at the time of labor market entry.

Intuitively, if we once again think of group a as males and group b as females, then $\phi_a > \phi_b$ for a variety of reasons. For example, entry level female employees may be less likely to be invited to informal social events such as golf outings or trips to sporting events by their male superiours, which may in turn give them fewer chances to make a good impression on the individuals who make promotion decisions. Similarly, different socialization across genders may result in females being less likely to speak up in meetings or conferences, once again giving females fewer chances to make positive impressions on potential employers or senior managers.[8]

Finally, assume that firms must contract with workers prior to performing each task, but workers have limited liability, meaning employers cannot make workers pay for any losses firms incur due to failed tasks. Hence, firms will not be able to tie wage payments to successes or failures. Furthermore, assume each worker knows his or her own type, chooses jobs and employers so as to maximize his or her wage, and can costlessly switch firms if a different employer offers a higher wage for a subsequent task. Finally, assume firms are expected profit maximizers and the labour market is competitive.

3 Equilibrium

A Perfect Bayesian Equilibrium (PBE) in this environment consists of: (i) a set of wage schedules and job assignment rules offered by each firm that constitute a Nash Equilibrium of a simultaneous move game given each firm's beliefs, and (ii) each firm's beliefs regarding each worker's type is calculated using Bayes' rule whenever possible (but can be anything for outcomes lying off the equilibrium path).

[8] There is a large disparate literature on such gender differences in communication and assertiveness. Particular examples include Babcock and Laschever (2003), Hollandsworth and Wall (1977), Kimble Marsh, and Kiska (1984), and the citations therein.

3.1 *Wages and Job Assignment*

To derive a PBE, first consider the wage schedule and job assignment rules. In general, for a given worker, firms must have the same equilibrium beliefs regarding his or her type since all firms have the same information. Therefore, given the beliefs about a certain worker's type, the unique Nash Equilibrium is to assign the worker to the job where he or she will have the highest expected revenue in the next task, and pay a wage equal to his or her expected revenue at that task.

To see why this behaviour constitutes a Nash Equilibrium, consider whether any firm has an incentive to deviate. Clearly, if a firm offers lower wage than prescribed by the proposed equilibrium, it will not be able to attract any workers. Similarly, if the firm offered a wage higher than the one prescribed above, it will lose money regardless of which job it assigns the worker to. Finally, it should be clear that if the firm offers a wage consistent with the proposed equilibrium, but assigns the worker to a different task than prescribed by the proposed equilibrium, the firm would also lose money. Hence, no firm has an incentive to deviate from the proposed equilibrium behaviour.

To see why this behaviour also constitutes the unique Nash equilibrium we must consider all other possible strategies. Since all firms are identical, we can restrict our analysis to only cases with symmetric strategies. For ease of reference, let $j^*(p)$ indicate the job where a worker with believed probability p of being of type-h has the highest expected productivity. Similarly, let $w^*(p)$ equal the expected productivity for this worker at a job $j^*(p)$ task. Moreover, note that in any Nash Equilibrium of this game, firms must be making zero expected profits or there is always an incentive for one firm to deviate and offer a slightly higher wage, which would mean it would capture all the workers, and therefore all the profits.

Given the above notation and the fact that all firms must earn zero expected profits in equilibrium, first consider a case where firms offer workers over whom they have beliefs p a wage w that is less than $w^*(p)$. Then, regardless of which job these firms assign these workers to, one firm can deviate by offering some w', where $w < w' < w^*(p)$, and assigning these workers to job $j^*(p)$. Since such a deviation would be a profitable deviation, it cannot be an equilibrium for firms to offer a wage less than $w^*(p)$, regardless of how they assign workers to jobs. Similarly, if firms offer a wage greater than $w^*(p)$, they will lose money regardless of how they assign workers to jobs. Therefore, given such strategies, each firm would have an incentive to deviate by shutting down (i.e. not hire any workers). Hence, it also cannot be an equilibrium for firms to offer a wage greater than $w^*(p)$, regardless of how they assign workers to jobs.

Finally, consider a case where firms pay workers over whom they have beliefs p a wage of $w^*(p)$, but assign them to a job $j \neq j^*(p)$. Since $w^*(p)$ is equal to this worker's highest possible expected productivity, assigning him or her to any other task besides $j^*(p)$ will clearly cause firms to lose money. Hence, each firm would once again have an incentive to deviate by shutting down, ruling out the possibility of any such equilibrium.

One thing to note about this equilibrium behaviour is that while this environment is dynamic in the sense that workers work for multiple periods and new information is revealed about each worker each period, equilibrium firm behaviour is essentially static. Specifically, each firm only takes into account expectations for what will happen in the subsequent period, not any periods further in the future. The reason for this is that workers can always leave a firm at any time and all firms observe the same information about each worker, regardless of who that worker actually works for. Therefore, firms do not have any incentive to promote workers as a sort of investment for obtaining better information about a worker's type. Specifically, if they were to do so, they would also reveal any information they observe to all other firms who did not incur any of the risks associated with promoting a worker, and who can freely hire the worker away if the incumbent firm doesn't offer that worker a wage equal to his or her updated expected productivity at his or her next task.

This equilibrium firm behaviour is quite easy to characterize with simple rules. Specifically, note that if a firm believes a worker has a probability p of being a type-h worker, the expected revenue from the subsequent task performed by this worker at job j equals

$$pR_j + (1-p)[(1-\pi_j)R_j - \pi_j L_j]. \tag{1}$$

Given this expected revenue, let p_j^* be the minimum probability such that the worker has a higher expected revenue for doing his or her next task at job j than job $j-1$, meaning p_j^* is the minimum p that solves

$$pR_j + (1-p)[(1-\pi_j)R_j - \pi_j L_j] \geq pR_{j-1} + (1-p)[(1-\pi_{j-1})R_{j-1} - \pi_{j-1}L_{j-1}],$$

or equivalently p_j^* is the minimum p that solves

$$\frac{p}{1-p} \geq \frac{\pi_j L_j - \pi_{j-1}L_{j-1}}{R_j - R_{j-1}} - \frac{(1-\pi_j)R_j - (1-\pi_{j-1})R_{j-1}}{R_j - R_{j-1}}.$$

Given the assumption that $L_j - L_{j-1} > R_j - R_{j-1}$, it can easily be confirmed that the above equation implies that $p_2^* > p_1^*$. Intuitively, this means p_1^* and p_2^* are threshold beliefs such that optimal firm behaviour is to promote/hire a worker to job 1 from job 0 if and only if it believes the probability that the worker is of type-h exceeds p_1^*, and similarly, promote/hire a worker to job 2 from job 1 if and only if it believes the probability that the worker is of type-h

exceeds p_2^*. Furthermore, if a firm ever comes to believe the probability that worker is of type-h is less than p_1^*, it should fire/demote the worker to job 0, regardless of which job that worker was previously assigned to.

Furthermore, given equation (1), the equilibrium wage offer to a worker who firms believe to be of type-h with probability p and who is optimally assigned to job j, will simply be given by

$$w_j^*(p) = pR_j + (1-p)[(1-\pi_j)R_j - \pi_j L_j]. \tag{2}$$

3.2 *Beliefs*

Recall that equilibrium beliefs must be determined using Bayes' rule whenever possible. To characterize these beliefs, first consider a firm's equilibrium beliefs regarding a member of group g, who has either not worked or only worked at job 0.[9] Let $p_0^g(n)$ denote the firm's beliefs for such a worker, given he or she has emitted n positive skill signals. For an individual who has yet to have any skill signalling opportunities (meaning $n = 0$), the firm has no individual specific information other than his group, implying that

$$p_0^g(0) = \alpha_g.$$

We can assume that $\alpha_g R_1 + (1-\alpha_g)[(1-\pi_1)R_1 - \pi_1 L_1] < R_0$ for $g = a, b$, implying $p_0^g(0) < p_1^*$, meaning firms never find it optimal to assign a worker to job 1 or job 2 when only the worker's group is observed. This means all workers must start out working at low-level jobs and have skill signalling opportunities arise each period with probability ϕ_g (for members of group g).

[9] Note that since everyone succeeds at job 0 tasks, no information is revealed about skill by observing a worker's performance at job 0.

Now take a group g worker who has only worked at job 0 or not worked at all, but who has had one skill signalling opportunity that resulted in a negative signal. Since only type-ℓ workers emit negative skill signals, firms will believe this worker to be of type-ℓ. More generally, regardless of skill signal history, once an individual emits a negative skill signal, firms will forever believe the worker to be of type-ℓ (since that failure is public information). Now consider a group g worker who has only worked at job 0 or not worked at all, but has had one opportunity to signal skill and emitted a positive signal. In this case, firms' equilibrium beliefs must be given by

$$p_0^g(1) = \frac{\alpha_g}{\alpha_g + (1-\alpha_g)(1-\lambda_g)},$$

which can be re-written as

$$p_0^g(1) = \frac{1}{1 + \frac{(1-\alpha_g)}{\alpha_g}(1-\lambda_g)}.$$

More generally, for a group g worker who has only worked at job 0 or not worked at all, and who has emitted n positive skill signals, firms' equilibrium beliefs must be given by

$$p_0^g(n) = \frac{1}{1 + \frac{(1-\alpha_g)}{\alpha_g}(1-\lambda_g)^n}, \tag{3}$$

which is increasing in n (given $0 < \lambda_g < 1$) and converges to 1 as n goes to infinity.

In order to determine equilibrium beliefs for workers employed at job 1, first let n_g^* denote the minimum number of positive skill signals an individual from group g must emit to be hired/promoted to job 1 (i.e. such that $p_0^g(n_g^*) = p_1^*$).[10]

[10] It is possible that there is no integer such that $p_0^g(n)$ equals p_1^* exactly. In this case, n_g^* can be considered the smallest integer for which $p_0^g(n) > p_1^*$, which will mean $p_0^g(n_g^*)$ will only be approximately equal to p_1^*. However, given n_g^* is relatively large, any difference will be

Note that since $p_1^* < p_2^*$, it will be true that n_g^* will be such that $p_0^g(n_g^*) < p_2^*$, meaning it will never be the case that individuals are promoted/hired for job 2 prior to starting at job 1. Given this, consider an individual with exactly n_g^* positive skill signals. By construction, firms will believe that the probability that this worker is of type-h equals $p_0^g(n_g^*) = p_1^*$ and therefore promotes/hires him to job 1.

Now, consider a worker who has worked at job 1 for one period, attempting one task. If he fails that task, the firm believes him to be of type-ℓ. Similarly, if a worker fails any subsequent task, then regardless of his/her task success or signalling history, the firm believes him/her to be of type-ℓ (causing the worker to go back to job 0 either due to demotion or firing). Alternatively, for a worker who has worked one period and completes his/her first task successfully, the firm believes him/her to be of type-h with probability $p_1(1)$, where

$$p_1(1) = \frac{1}{1 + \frac{(1-p_1^*)}{p_1^*}(1-\pi_1)}.$$

Note that the above expression does not need a superscript g since it does not depend on a worker's group. This means that once promoted to job 1, the *only* thing that matters for firm beliefs is whether the worker successfully completes his/her tasks. *The fact that the worker was promoted is a sufficient statistic regarding firms' beliefs with respect to a worker's type at the beginning of his or her job 1 career.* More generally, after a worker successfully completes s_1 job 1 tasks, the firm believes him/her to be of type-h with probability $p_1(s_1)$, where

$$p_1(s_1) = \frac{1}{1 + \frac{(1-p_1^*)}{p_1^*}(1-\pi_1)^{s_1}}.$$

Equilibrium beliefs regarding workers at job 2 can be constructed similarly

negligible.

to beliefs regarding workers at job 1. In particular, let s^* denote the minimum number of successful job 1 tasks such that $p_1(s^*) = p_2^*$. Essentially, this means s^* can be interpreted as the observable qualifications necessary for promotion up to job 2, or the director level job. Note that this qualification standard, or successful task completion threshold, does not depend on a worker's group g. This is because, as discussed above, by virtue of being promoted/hired for job 1, workers essentially start from scratch, all with the same believed probability of being a type-h worker regardless of group or skill signalling history (namely p_1^*).

A worker with s^* successful job 1 task completions will be offered a promotion and will take it since the equilibrium wage (given the implicit firm beliefs about such a worker) will dominate what he or she is being paid for tasks at job 1. Therefore, a firm's beliefs regarding a worker who has just been promoted to job 2 but has yet to complete any tasks will equal $p_1(s^*) = p_2^*$. Similar to above, if a worker fails at any job 2 tasks, the firm will believe to be of type-ℓ and will demote the worker back to job 0 (or equivalently fire the worker). Alternatively, after s_2 successful task completions at job 2, firms' beliefs regarding the probability a worker is of type-h will equal

$$p_2(s_2) = \frac{1}{1 + \frac{(1-p_2^*)}{p_2^*}(1-\pi_2)^{s_2}}.$$

Finally, assume that firms believe any worker who turns down a promotion, or chooses to take a voluntary demotion, is of type-ℓ. Such beliefs are not only reasonable, as a type-h worker always has the incentive to accept a promotion and stay at the highest job level possible since he will earn higher wages and more easily reveal his skill (since $\lambda_g < \pi_1 \leq \pi_2$), but such beliefs will also ensure that no one ever turns down a promotion or volunteers to take a demotion. Given no one ever turns down a promotion or volunteers to take a demotion in equilibrium, these beliefs are for occurrences that lie off the equilibrium path and therefore are not restricted in a Perfect Bayesian Equilibrium.

3.3 *Analyzing the Equilibrium*

The equilibrium derived above essentially works as follows. Workers from group g start out working at the low level job (which can also be interpreted as school), during which they are paid

$$w_0 = R_0 \tag{4}$$

and have a probability ϕ_g of getting a skill signalling opportunity each period. If they ever emit a negative skill signal, they work at the low level job for the rest of their career. Alternatively, once they emit n_g^* positive skill signals, they are hired/promoted to job 1. Once at job 1, they are fired/demoted to job 0 if they ever fail a job 1 task, but are paid

$$w_1^*(p_1(s_1)) = p_1(s_1)R_1 + (1 - p_1(s_1))[(1 - \pi_1)R_1 - \pi_1 L_1] \tag{5}$$

prior to the next task they are assigned, given they have already successfully completed s_1 job 1 tasks. Finally, if they complete s^* job 1 tasks successfully before they reach retirement age, the worker is hired/promoted to job 2, where once again the worker is fired/demoted to job 0 if he or she ever fails a job 2 task, but is paid

$$w_2^*(p_2(s_2)) = p_2(s_2)R_2 + (1 - p_2(s_2))[(1 - \pi_2)R_2 - \pi_2 L_2] \tag{6}$$

prior to the next task he or she is assigned, given he or she has already successfully completed s_2 job 2 tasks.

There are several interesting implications resulting from this equilibrium.

First, even if $\alpha_a > \alpha_b$, $\lambda_a > \lambda_b$, and/or $\phi_a > \phi_b$, there is *no wage discrimination in equilibrium* (even though firms are not restricted in any way from doing so!), as can be seen in equations (4)-(6) and noting that worker's group g does not appear in these equations. Once working at any particular job, wages are only a function of the successfully completed tasks at that job, not the worker's group or previous skill signalling or task success history at a different job. Second, there is also *no discrimination with respect to hiring/promotion to job 2*, or the director level. A worker simply has to successfully complete s^* job 1 tasks in order to be deemed "qualified" for job 2, regardless of the worker's group affiliation.[11]

Discrimination only arises in one aspect of the labour market in equilibrium, namely in hiring/promotion to job 1, or entering the career-track jobs. From equation (3) we can see that if $\alpha_a > \alpha_b$ and/or $\lambda_a > \lambda_b$, then $p_0^a(n) > p_0^b(n)$ for any given n. Therefore, we know that if $\alpha_a > \alpha_b$ and/or $\lambda_a > \lambda_b$ it must be the case that $n_a^* < n_b^*$, or that members of group b have to emit more positive skill signals in order to convince employers to hire them for career-track jobs than members of group a.[12]

These results are interesting in that the discrimination that arises in this model is arguably the most difficult to combat. Specifically, if individuals from

[11] This model of promotion bears some resemblance to Meyer (1991), in that firms must make promotion decisions based on observing each individual's track record for successfully completed tasks. However, in Meyer's model, a firm is trying to use its observed information to determine which of two workers is more productive, in order to promote the more productive worker. In the model presented here, a firm can promote as many workers as it wants, however, it only wants to promote those workers who are of the right type to do the job successfully. Hence, the optimal promotion strategy in this model differs qualitatively from that in Meyer's model.

[12] Note that these implications are very close in spirit to those coming from Lazear and Rosen (1990) in the context of gender inequality. Namely, in their model, they show that males and females will be paid similarly within the same job. However, females face discrimination with respect to promotion, with differences in promotion rates by gender being smaller at very high levels of ability than at middle or low levels. As discussed in the introduction above, however, Lazear and Rosen's results arise due to a very different underlying mechanism. Rather than differences in average skill level, skill signalling precision, and/or skill signalling opportunities across genders, their key assumption is that females on average value time at home more than similarly skilled men.

different groups with similar successes in their career-track jobs were being paid differently, or had different standards for promotion to the highest level director jobs, discrimination lawsuits should be relatively easy to win as workers can simply show that they were treated differently than workers of another group with the same job success history. However, the "skill signals" relevant for hiring/promotion to the career-track jobs might be more subjective, and this is where discrimination is occurring.

For example, suppose relatively fewer females than males become highly qualified in science (i.e. α is greater for males than females with respect to science) as suggested by former Harvard President Lawrence Summers (Summers, 2004).[13] Then, according to this model, when it comes to hiring for entry level research jobs in the sciences, employers will require more positive skill signals from females than males for a given position. This could take the form of requiring females to have better or more letters of reference and stronger interviews than males. While such practices may be discriminatory, they are difficult to police or legislate against, and are arguably quite prevalent. However, once hired onto the track of research scientist, highly skilled females are held to the same standards as their male colleagues when it comes to further promotions.

4 The "Glass Ceiling"

Note that the results arising from the equilibrium derived above are consistent with the notion that the reason one group is underrepresented at the highest level jobs is not because of discrimination with respect to promotion to these jobs, but rather because there are not enough "qualified applicants". However, in the context of this model, "qualified applicants" does not mean those who can

[13] As Summers points out, such a relative scarcity of females at the top of end of the skill distribution may be due to either a difference in mean skill level between males and females, or simply a difference in the variance of skill between males and females.

successfully do the job (i.e. type-h workers). Rather, it means those who have developed the requisite success history during their time in job 1 that allows them to be promoted to job 2. However, as will be shown in detail below, just because there is no discrimination with respect to promotion to director level jobs, it will not be the case that equally skilled workers from different groups have the same probability of eventually being promoted to the director level. Indeed, even if the two groups have the same skill distribution (i.e. $\alpha_a = \alpha_b$), group b workers will still be underrepresented at the job 2 level as long as $\lambda_a > \lambda_b$ and/or $\phi_a > \phi_b$ (with this underrepresentation becoming even more severe if it is also true that $\alpha_a > \alpha_b$).

We can calculate the fraction of each type from each group that makes it to job 2 as follows. In order to make it to job 2 during one's career, a worker needs to work at job 1 for s^* periods in order to get the requisite number of job 1 successes to be promoted. Given the worker can only work for T periods in his or her entire work life, the worker has to be hired/promoted to job 1 in less than $T - s^*$ periods in order to make it to job 2. Therefore, in order to make it to job 2, a worker from group g has to emit n_g^* positive skill signals in less than $T - s^*$ periods. Noting that the probability of having a skill signalling opportunity each period for a member of group g is ϕ_g, the fraction of type-h individuals from group g who will emit the required number of positive skill signals in the required time period (denoted F_g^h) is given by the following Binomial distribution

$$F_g^h = \sum_{k=n_g^*}^{T-s^*} \binom{T-s^*}{k} \phi_g^k (1-\phi_g)^{T-s^*-k}. \tag{7}$$

Similarly, the fraction of type-ℓ individuals from group g who will emit the required number of positive skill signals in the required time period (denoted F_g^ℓ) is given by

$$F_g^\ell = \sum_{k=n_g^*}^{T-s^*} \binom{T-s^*}{k} \phi_g^k (1-\phi_g)^{T-s^*-k} (1-\lambda_g)^{n_g^*} (1-\pi_1)^{s^*}. \quad (8)$$

Equations (7) and (8) not surprisingly show that a higher fraction of type-h workers will make it to job 2 than type-ℓ workers (since $0 < \lambda_g < 1$ for both groups g). Furthermore, they also lead to the following key proposition.

Proposition 1 *Individuals from group b will be less likely than equally skilled individuals from group a to make it to director level jobs (i.e. job 2) if one or more of the following conditions hold:*

1. Group a is more skilled on average than group b (i.e. $\alpha_a > \alpha_b$),
2. The skill signals are more precise for group a than group b (i.e. $\lambda_a > \lambda_b$),
3. Members of group a generally have more frequent skill signalling opportunities than members of group b (i.e. $\phi_a > \phi_b$).

Proof. Direct from equations (7) and (8), and recalling that $n_a^* > n_b^*$ if $\alpha_a > \alpha_b$ and/or $\lambda_a > \lambda_b$. ■

It is relatively obvious that if one group has a higher fraction high skilled members than another group, the more highly skilled group will have a higher fraction of its members make it to the highest job level. However, Proposition 1 is much stronger than this. It says that if one group has a higher fraction of high skilled members than another group, then members of the more skilled group are more likely to make it to the highest job level than *equally skilled* members of the other group. Hence, even while there is no direct discrimination with respect to being promoted to the highest job level (i.e. job 2), there is an inherent inequality of opportunity with respect to getting there across the two groups. Proposition 1 goes on to say that this inherent inequality of opportunity will also arise, or be further exacerbated, if pre-market or early job skill signals are less precise for one group than the other and/or if one group generally has more opportunities to signal their skill early in their life than the other group.

To summarize, the above analysis shows how a "glass ceiling" phenomenon may arise not because members of group b face greater hurdles in order to be promoted from mid-level jobs to top jobs than members of group a, but rather because members of group b face a "sticky floor" in the sense that they face greater hurdles than group a with respect to being promoted out of the entry level jobs and into mid-level career-track jobs. Such delays early on may cause many group b workers to not have sufficient time to develop the success record at mid-level jobs required for promotion to the top jobs in the economy.

5 How Does This Model Relate other Group Distinctions Besides Gender?

The model developed above was generally motivated by thinking about how inequality of opportunity with respect to hiring and promotion may arise between gender groups. In this context, one further issue that was not discussed was how career breaks for maternity leave may affect females. As the above model makes clear, if females are more likely than males to take time out of the labor market, they will have fewer opportunities to signal their skill and/or fewer opportunities to successfully complete tasks, both of which will impede their chances of making it to the top jobs in their professions relative to males.[14]

Unlike many of the other models that examine gender inequality in the labour market however, this model can be straightforwardly applied to a variety

[14] For further discussion of why females may be more likely to leave the labor market than otherwise similar males, see Francois (1998) and Bjerk and Han (2007).

of other group distinctions as well. For example, groups a and b can be interpreted as referring to whites and blacks respectively. Due to systematic quality differences in the schools attended by whites and blacks (Card and Krueger, 1992), it is clear that α for whites is likely higher than it is for blacks. Moreover, since blacks often attend different schools than whites, white employers may find it harder to evaluate the relative value of things like grades, class rank, honor societies, and letters of recommendation or references for black applicants, causing λ to be greater for whites than blacks. Finally, the greater resources in primarily white schools relative to primarily black schools mean whites may have more opportunities to participate in extra-curricular activities and develop closer relationships with teachers. Such differences mean blacks may have fewer opportunities to signal their skill, or a lower ϕ, relative to whites. As discussed in Proposition 1, such differences across races would cause blacks to be less likely to make it to the top jobs in the economy than equally skilled whites.

Another application of this model is to think about underrepresentation of immigrants versus native born workers in the top job levels. Once again, due to schooling quality differences across countries, it could be true that immigrants have a relatively smaller fraction of individuals with the requisite skill level to succeed at these top job levels (i.e. α is lower for immigrants than native born workers). However, it is also very likely true that skill signals from immigrants are much less precise than those for native born workers (i.e. immigrants have a lower λ than native borns). For example, how should a British employer interpret graduating with honours from Calcutta University relative to graduating with honours from Cambridge? How should one interpret standardized verbal test scores from a non-native English speaker relative to similar scores for native English speakers? As shown in the model above, such greater skill signal imprecision can lead to immigrants being underrepresented at top job levels, even if immigrants have a very similar fraction of highly skilled workers as native borns.

6 How Does This Model Relate to the Empirical Findings Regarding Labour Market Inequality Between Groups?

The primary aim of the model developed above was to illustrate how it is possible for members of one group to have lesser opportunities for reaching the highest jobs in the economy than equally skilled members of another group, even if there is no discrimination with respect to promotion to these top jobs and there is no wage discrimination. Given this aim, the model was intentionally quite simple and therefore abstracted a great deal from the operation of labor markets in the real world. Therefore, it is not necessarily straightforward to take this model directly to the data.

However, even with the above caveat, it is still informative to consider the key implications of this model in light of the current empirical literature on labour market inequality between identifiable groups. Namely, is there evidence that earnings differences between distinct groups arise primarily due to inter-group differences in hiring and promotion rather than inter-group differences in pay within job level? Moreover, is there any evidence to suggest that such inter-group differentials hiring or promotion are generally relegated to lower level jobs, not the more lucrative high level jobs?

The evidence with respect to discrimination between equally skilled members of different groups in hiring is quite strong. With respect to gender, one key example is Goldin and Rouse (2000), who present convincing evidence of historical discrimination against females in major American Orchestras. Specifically, they look at the extent to which moving to blind auditions, where judges could not observe the gender of the musician, affected the likelihood that females were hired. They find that moving to blind auditions led to dramatic increases in the likelihood that females advanced to later audition rounds during the hiring process, and increased the likelihood a female was actually hired by 25 percent. Furthermore, this finding is arguably consistent with the statistical discrimination model derived above, as not only do Goldin and Rouse quote several major conductors as saying that they *believe* women are generally inferior to men in an orchestra, but also, the female musicians in their sample appear to be somewhat less “skilled” on average than the male musicians.[15]

In the context of race, similar evidence of discrimination in hiring comes from Bertrand and Mullainathan (2004). In this study, they sent resumes to a variety of newspaper job postings, where the only meaningful difference across resumes was that some had typically African-American names while others had typically white names. They found that the resumes with typically white names garnered roughly 50 percent more call-backs than those with typically African-American names.

With respect to promotion and within job wage inequality, the empirical literature on has generally focussed on differences across genders. Studies examining such issues find similar patterns exist across a variety of job sectors. For example, as far back as 1973, Makiel and Makiel found that in one large corporation, women tended to be paid similar salaries as men in the same job level, but generally earned less overall due to their lower average job levels. Subsequent work by Cabrel, Ferber, and Green (1981) found that females re-

[15]This result is suggested by the fact that including individual fixed effects are crucial for Goldin and Rouse’s results. When such effects are not included, females were found to be significantly less likely to advance to the next round of auditions than males even when auditions were blind.

ceived lower initial job placements and were less likely to be promoted than men with similar credentials at three large banking firms. Similarly, in examining the promotion rates of managers in a large Canadian corporation, Cummings (1988) found that females were "distinctly less likely than their male colleagues to be promoted, and, furthermore, that their disadvantage is not primarily the result of differential 'returns' to particular acquired attributes." Analogous results were found for the legal profession throughout the 1970s (Spurr, 1990). In analyzing a nationally representative sample of workers, Olson and Becker (1983) also conclude that "a woman is significantly less likely than a man to receive a promotion, other things equal" and that "unequal access to opportunities, rather than unequal returns, constitutes the principal source of male-female differences in employment outcomes."

In more recent studies, Winter-Ebmer and Zweimuller (1997) find that in the Austrian labour market data they analyze, "neither the risk of childbearing nor different productive characteristics can explain the crowding of females into lower hierarchical positions. Females have to fulfill higher standards to be promoted" (out of the lower hierarchical positions). In the context of academics, Ginther and Hayes (2003) look at data from the Survey of Doctorate Recipients and conclude that "gender discrimination for academics in the humanities tends to operate through differences in promotion, which in turn affects wages." Further evidence of differential promotion patterns across genders, but similar pay within job category, can be found in Blau (1977), Groshen (1991), Petersen and Morgan (1995), and Bertrand and Hallock (2002).

While the above studies suggest gender discrimination primarily arises in hiring and promotion, rather than differential wages within a given job, they do not necessarily speak to the other key implication that is unique to the model presented above—namely that any discrimination in hiring and promotion should generally be relegated to the lower end of job ladders. However, there also exists empirical evidence in support of this implication. For example, Jones and Makepeace (1996) examine gender differences in promotion in a large British financial institution and find little differences in salary between men and women at any particular job grade, but that "differential barriers (to promotion) confronting women are greatest at the lower end of the job ladder." However, they go on to conclude that men and women appear to "receive equal treatment once senior job grades have been reached."

More direct evidence related to the model developed above comes from Peterson and Saporta (2004). Using personnel records for managerial, administrative, and professional employees hired into a large U.S. firm between 1978 and 1986, they find evidence that among workers who start at the lowest job level, males are promoted to the next level at significantly faster rate than females of similar age and education level. However, they find no evidence that men are promoted at a faster rate than similar females for any job level thereafter. Moreover, at the time of hire, males are hired into somewhat higher job levels than females of similar age and education level, even among workers in their mid-twenties who generally have relatively little job experience. Peterson and Saporta estimate that it takes about four years of seniority for females to make up for these disadvantages relative to males at the outset of their careers with this firm.

Finally, in an analysis of U.S. government workers, where a promotion is clearly indicated by an increase in salary grade, DiPrete and Soule (1988) find that "women had difficulty advancing in the lower and middle grades" but "women and men had similar promotion rates in the higher grades, though women were underrepresented at these levels."

While the above studies provide a substantial amount of evidence consistent with the model of promotion developed in this paper, all these results are subject to the caveat that employers generally have more information regarding worker productivity than is available to the analyst. Therefore, while the male-female differences in promotion could be due to discrimination, it is also possible that they reflect unmeasured productivity differences between males and females. The findings showing that the gender differences in promotion rates are more pronounced at lower rungs of the job ladder mitigates this concern somewhat, but by no means overcomes it.

Moreover, the model above specifically suggests that similar differential rates of hiring or promotion out of entry level positions should be found between other distinguishable groups besides males and females, such as blacks and whites or immigrants and natives. While there is a very large literature on earnings inequality between blacks versus whites and immigrants versus natives, to my knowledge there has been no empirical work specifically analyzing promotion differences between such groups. The model developed above suggests that such analyses constitute an important next step toward furthering our understanding earnings inequality between such groups.

7 Policy Considerations

The inequality of opportunity with respect to making it to the top job level that arises in the model developed above can arguably provide an equity rationale for government intervention. However, standard affirmative action type policies will not be without costs. To see why, say a policy was implemented such that members of the underrepresented group needed fewer job 1 successes in order to be promoted to job 2 than members of the other group. Such a policy would certainly increase the fraction of the underrepresented group making it to the highest job level. However, there would be efficiency costs associated with such a policy because, unless firms start wage discriminating, firms would on average lose money on those members of the underrepresented group who were promoted to the top job. This may make firms more hesitant about hiring members of the underrepresented group in the lower level jobs. Moreover, members of the underrepresented group who are promoted to the highest job level will be believed (by firms and other workers) to be less skilled on average than members from the other group who make it to the top job level (even those that are not), which may cause tension within the workplace.

On the other hand, a policy could be implemented that attempts to decrease the number positive skill signals necessary for hiring/promotion to job 1 for members of the underrepresented group. Since inequality across groups with respect to the necessary number of positive skill signals required for hiring/promotion to job 1 is how discrimination manifests itself in this model, such a policy seems potentially superior. However, as alluded to above, implementing such a policy may be quite difficult, as the subjectivity of skill signals may make it hard to enforce such a policy. Moreover, even if such a policy could be implemented and enforced, it will have similar costs to the policy discussed in the previous paragraph. Specifically, unless firms start wage discriminating, firms will lose money on average when promoting members of the underrepresented group using the lower positive skill signal threshold as directed by the policy, possibly causing firms to shy away from recruiting and interviewing members of the underrepresented group in the first place. Moreover, such a policy will cause firms to start discriminating with respect to promotion to job 2 (i.e. it will cause firms to raise the number of job 1 successes necessary for promotion to job 2 for members of the underrepresented group). Also, like above, members of the underrepresented group will be believed to be less skilled on average than members of the other group at similar job levels.

While the above discussion does not necessarily imply that equity gains associated with either of the affirmative action type policies discussed above are outweighed by efficiency losses (as indeed any such calculation would be highly speculative), it does reveal that such affirmative action policies conform to the "no free-lunch" principle. Therefore, if possible, it will generally be preferable to affect the underlying parameters that actually lead to the inequality in the first place. Namely, optimal policies are likely policies that diminish differences between groups with respect to their skill distributions, the precision with which individuals can signal their skills to employers, and the frequencies with which individuals can signal their skill to employers prior to working or while working at low-level jobs. Policies that potentially may have many of these attributes are policies that cause employers to target recruiting and mentoring efforts toward underrepresented groups. Specifically, if targeted recruiting increases the number of individuals from the underrepresented groups who are interviewed, employers may become better at interpreting skill signals from these groups (i.e. increase λ for these groups), and would give members of these underrepresented groups more opportunities to signal their skill (i.e. increase ϕ for these groups). Similarly, targeting mentoring programs at underrepresented groups may not only influence human capital decisions (possibly increasing α for the underrepresented groups), but may also increase the skill signalling opportunities for underrepresented groups (i.e. increase ϕ).

8 Summary and Conclusions

How productive a given worker will be at a particular job is not generally known by employers with certainty, especially at higher level jobs that require a variety of tasks and unforeseen challenges. Rather, firms must base their hiring and promotion decisions on expectations of how productive a worker will be. This paper shows how modelling this process in a dynamic setting, where firms accumulate increasing amounts of information about each worker over the course of his or her career, can provide some new insights into the issue of the glass ceiling phenomenon, where females and minorities are underrepresented at the top job levels even though there does not appear to be explicit discrimination with respect to promotion to these top jobs.

In particular, the model shows how three distinct factors may cause equally skilled members from different groups to have different expected labour market outcomes, even when firms do not inherently "prefer" one group over the other and firms observe a relatively large number of informative skill signals about each worker over the course of his or her career. Specifically, inequality of opportunity will arise if (i) one group has a smaller fraction of highly skilled individuals than another, (ii) skill signals emitted by members of one group are evaluated less precisely by those hiring than skill signals from anther group, and (iii) one group generally has less frequent opportunities to signal their skill prior to starting their career or early in their career than do members of another group.

All three of the above possible differences between groups can result in longer periods of time and effort required for members of one group to be hired/promoted to career-track jobs than for equally skilled members of another group. The first two possible group differences lead to such a result because such differences will cause firms to require more positive skill signals from one group than the other in order to be sufficiently convinced of their skill level to hire/promote them to career-track jobs. The third factor leads to this result because it will mean that it will generally take members of one group longer to accumulate the number of positive skill signals required for hiring/promotion to the career-track job.

Such delays in getting to a career-track job level for members of one group relative to members of another will then result in relatively fewer individuals from one group being able to put the success record together while at the career-track job necessary to qualify themselves for hiring/promotion to the highest level jobs in the relevant time frame. Therefore, while the model shows that firms will not discriminate with respect to hiring/promotion to the highest job levels in the economy (or with respect to wages within a job level), there may still be inequality of opportunity. Namely, individuals from one group may experience delays in the early parts of their careers due to the reasons discussed above, causing them to be less likely to eventually make it to the highest level jobs than similarly skilled members of another group.

This potential for inequality of opportunity with respect to making it to the top job level can justifiably be a reason for concern if society values principles of equality opportunity across groups. However, as discussed above, affirmative action policies meant to remedy inequality of opportunity are likely to have both efficiency costs and/or may lead to discrimination in other aspects of the labour market. If possible, the model indicates that it would be preferable to implement policies that eradicate differences in human capital attainment, skill signalling precision, and skill signalling opportunities between groups. Certainly more research is necessary to determine whether there exist feasible policies that can have such effects, however, policies encouraging targeted recruiting initiatives and mentoring programs for underrepresented groups are likely to be steps in the right direction.

Claremont McKenna College

References

Aiger, D.J. and Cain, G. (1977). "Statistical theories of discrimination in the labor market." *Industrial and Labor Relations Review,* vol. 30, pp. 175-187.

Albrecht, J., Bjorklund, A., and Vroman, S. (2003). "Is there a glass ceiling in sweden?" *Journal of Labor Economics*, vol. 21, pp. 145-177.

Athey, S., Avery, C., and Zemsky, P (2000). "Mentoring and diversity." *American Economic Review,* vol. 90, pp. 765-786.

Babcock, L. and Laschever, S. (2003). *Women Don't Ask: Negotiation and the Gender Divide*. Princeton, NJ: Princeton University Press.

Becker, G.S. (1981). *A Treatise on the Family*. Cambridge, MA: Harvard University Press.

Becker, G.S. and Tomes, N (1979). "An equilibrium theory of distribution of income and intergenerational mobility," *Journal of Political Economy,* vol. 87, pp. 1153-1189.

Bertrand, M. and Hallock, K.F. (2002). "The gender gap In top corporate jobs," *International Labor Relations Review,* vol. 55, pp. 3-21.

Bertrand, M. and Mullainathan, S. (2004). "Are Emily and Brendan more employable than Lakisha and Jamal? A field experiment on labor market discrimination," *American Economic Review*, vol. 94, pp. 991-1013.

Bjerk, D. and Han, S. (2007). "Assortative marriage and the effects of government homecare provision programs on gender wage and participation inequality." *Journal of Public Economics,* vol. 91, pp. 1135-1150.

Cabral, R., Ferber, M.A., and Green, C.A. (1981). "Men and women in fiduciary institutions: a study of sex differences in career development." *Review of Economics and Statistics*, vol. 63, pp. 573-80.

Cannings, K. (1988). "Managerial promotion: the effects of socialization, specialization, and gender." *Industrial and Labor Relations Review*, vol. 42, pp. 77-88.

Card, D. and Krueger, A. (1992). "School quality and black-white relative earnings: a direct assessment." *Quarterly Journal of Economics*, vol. 107, pp. 151-200.

Cornell, B. and Welch, I. (1996). "Culture, information and screening discrimination," *Journal of Political Economy*, vol. 104, pp. 542-571.

Coate, S. and Loury, G. (1993). "Will affirmative action eliminate negative stereotypes?" *American Economic Review*, vol. 83, pp. 1220-40.

Committee on the Status of Women in the Economics Profession (CSWEP). (1999). *1999 Annual Report*.

DiPrete, T.A. and Soule, W.T. (1988). "Gender and promotion in segmented job ladder systems." *American Sociological Review*, vol. 53, pp. 26-40.

Equal Opportunities Commission (EOC). (2005). *Sex and Power: Who Runs Britain? 2005.*

Federal Glass Ceiling Commission. (1995). *A Solid Investment: Making Full Use of the Nation's Human Capital.* Washington D.C. November.

Francois, P. (1998). "Gender discrimination without gender difference: theory and policy responses," *Journal of Public Economics*, vo. 68, pp. 1-32.

Fryer, R. (forthcoming). "Belief flipping in a dynamic model of statistical discrimination," *Journal of Public Economics*.

Ginther, D. and Hayes, K. (2003). "Gender differences in salary and promotion for faculty in the humanities 1977-95," *Journal of Human Resources*, vol. 38, pp. 34-73.

Goldin, C. and Rouse. C. (2000). "Orchestrating impartiality: the impact of 'blind' auditions on female musicians," *American Economic Review*, vol. 90, pp. 715-741.

Hollandsworth, J. G. and Wall, K.E. . (1977). "Sex differences in assertive behavior: An empirical investigation," *Journal of Counseling Psychology*, vo. 24, pp. 217-222.

Jones, D. and Makepeace, G. (1996). "Equal worth, equal opportunities: pay and promotion in an internal labour market," **ECONOMIC JOURNAL**, vol. 106, pp. 401-409.

Kimble, C. E., Marsh, N. B., and Kiska, A. C. (1984). "Sex, age, and cultural differences in self-reported assertiveness." *Psychological Reports*, vol. 55, pp. 419-422.

Lang, K. (1986). "A language theory of discrimination," *Quarterly Journal of Economics*, vol. 101, pp. 363-382.

Lazear, E. and Rosen, S. (1990). "Male-female wage differentials in job ladders," *Journal of Labor Economics*, vol. 8, pp. S106-S123.

Lundberg, S. and Startz, R. (1983). "Private discrimination and social intervention in competitive labor markets," *American Economic Review*, vol. 73, pp. 340-347.

Malkeil, B.G. and Malkiel, J.A. (1973). "Male-female pay differentials in professional employment," *American Economic Review,* vol. 63, pp. 693-705.

McDowell, J., Singell, L., and Ziliak, J. (1999). "Cracks in the glass ceiling: gender and promotion in the economics profession." *American Economic Review, Papers and Proceedings of the 111 Annual Meeting*, vol 89, pp. 392-396.

Meyer, M. (1991). "Learning from coarse information: biased contests and career profiles," *Review of Economic Studies,* vol. 58, pp. 15-41.

Milgrom, P. and Oster, S. (1987). "Job Discrimination, Market Forces, and the Invisibility Hypothesis." Quarterly Journal of Economics 102: 453-476.

Moro, A., and Norman, P. (2004). "A general equilibrium model of statistical discrimination," *Journal of Economic Theory*, vol. 114, pp. 1-30.

Olson, C.A. and Becker, B.E. (1983). "Sex discrimination in the promotion process," *Industrial and Labor Relations Review*, vol. 36, pp. 624-641.

Petersen, T. and Saporta, I. (2004). "The opportunity structure for discrimination," *American Journal of Sociology*, vol. 109, pp. 852-901.

Phelps, E.S. (1972) "The statistical theory of racism and sexism," *American Economic Review*, vol. 62, pp. 659-61.

Rosenberg, M. (2004). "Two Few–of 117 football coaches, two are black; its called institutional racism." *Detroit Free Press* (December 2, 2004).

Spurr, S.J. (1990). "Sex discrimination in the legal profession: a study of promotion," *Industrial and Labor Relations Review*, vol. 43, pp. 406-417.

Summers, L. (2005) "Remarks at NBER conference on diversifying the science & engineering workforce," http://www.president.harvard.edu/speeches/2005/nber.html (accessed July, 2005).

Tannen, D. (2001). *You Just Don't Understand: Men and Women in Conversations*. New York: Harper Collins.

Winter-Ebmer, R. and Zweimuller, J. (1997). "Unequal assignments and unequal promotion in job ladders," *Journal of Labor Economics*, vol. 15, pp. 43-71.