Second-Order Statistical Discrimination*

Tilman Klumpp[†]
University of Alberta

Xuejuan Su[‡]
University of Alberta

August 2012

Abstract

The low representation of female workers in elite jobs is sometimes attributed to a tail effect: If the human capital distribution exhibits less variation among females than among males, then even with comparable average human capital there will be fewer females in the right tail than males. This paper offers an explanation for why the human capital distribution might have this property. We show that the belief that the female human capital distribution has a lower variance than the male distribution can be self-fulfilling, in that it provides individuals with incentives to invest in human capital such that the resulting distribution exhibits exactly this characteristic. If this happens, fewer females are employed in high-end jobs (a "glass ceiling" effect). The average productivity of female workers may at the same time be higher than that of male workers.

Keywords: Statistical discrimination; aptitude distribution; gender differences; "glass ceilings".

JEL codes: J71, J78.

*We thank Maria Arbatskaya, Paul Peccorino, Paul Rubin, Akram Temimi, the editor Hanming Fang, and two anonymous referees for helpful comments.

[†]Department of Economics, University of Alberta. 9-20 H.M. Tory Building, Edmonton, AB, Canada T6G 2H4 (klumpp@ualberta.ca).

[‡]Department of Economics, University of Alberta. 9-22 H.M. Tory Building, Edmonton, AB, Canada T6G 2H4 (xuejuan1@ualberta.ca).

1 Introduction

Job applicants may be treated differently by employers based on their race, gender, or other group membership—even if they are similar in characteristics that are more directly related to their skills, such as their grades. This behavior is rational for employers if an individual's true qualification for a job cannot be perfectly observed *and* groups differ in the average qualification of their members. In this situation, group membership conveys statistical information about an individual's qualifications, and economists use the term *statistical discrimination* to describe the resulting differential treatment of individuals based on the group they belong to. In seminal contributions, Arrow (1973), and later Coate and Loury (1993), demonstrate that the inter-group differences that cause statistical discrimination can also be the result of statistical discrimination: Two *ex ante* identical groups can end up with different average skills if the prospect of discrimination in the labor market discourages one group from investing in their skills. In other words, differences in average qualifications across groups can become self-fulfilling expectations—an idea which has had a profound impact on economists' understanding of labor market inequalities.[1]

In this paper, we extend the self-fulfilling expectations model of statistical discrimination from the first to the second moment of the skill distribution. We show that the same mechanism that can explain differences in means across populations can also explain differences in variance—a phenomenon we call *second-order statistical discrimination.* To motivate this exercise, consider the question why fewer women than men are employed in high-end jobs in business, science, or engineering. One hypothesis is that women are less able than men in these fields, on average. A second hypothesis holds that men and women are similarly able on average, but that the distribution of ability has a higher variance among men than among women. Fewer women than men are then located in the right tail of the ability distribution, which is the relevant region for high-end jobs.[2] This hypothesis raises the question: *Why* would the ability distribution be more variable among men than among women? The notion of second-order statistical discrimination explored in this paper offers a possible explanation.[3]

[1]Most importantly, in order to reduce labor market inequalities arising from (self-fulfilling) statistical discrimination, different policies are required than if discrimination is merely taste-based. While free labor markets tend to self-correct toward a non-discriminatory state in the latter case (see Becker 1957), government intervention such as affirmative action policies can be required in case of statistical discrimination. For a thorough survey of the literature of statistical discrimination and policy responses to it, see Fang and Moro (2011).

[2]The dispersion hypothesis was famously put into the spotlight when former Harvard president Larry Summers made the following remarks at an NBER conference, for which he was later criticized by some colleagues: "It does appear that on many different human attributes [...] there is relatively clear evidence that, whatever the difference in means, which can be debated, there is a difference in the variability of a male and a female population. [...] If one is talking about people who are three-and-a-half, four standard deviations above the mean, even small differences in standard deviation will translate into very large differences in the available pool substantially out" (Summers 2005).

[3]There also exists an alternative, biological explanation of second-moment differences in the male and female human capital distribution (see Rubin and Paul 1979; Browne 1998; Browne 2006, and references

Our model can be summarized as follows. An individual's final ability, or human capital, is the product of his or her innate ability as well as a costly personal investment. Neither an individual's innate ability, nor the investment, nor the resulting final human capital are directly observable. Suppose, now, that innate ability follows the same distribution in both men and women, but that the individuals' investment decisions result in a human capital distribution that has a higher variance among men than among women. If employers observe an imperfect signal of each individual's human capital, then a low signal from a male applicant must indicate relatively low human capital, compared to the same signal from a female applicant. Similarly, a high signal from a male applicant indicates high human capital, compared to the same signal from a female applicant. In assigning workers to jobs, employers will therefore discriminate against men with relatively low signals, and at the same time against women with relatively high signals. The first effect results in a "sticky floor" for men, and the second in a "glass ceiling" for women. We show that this pattern of discrimination discourages women of high ability, as well as men of low ability, from investing into human capital. The resulting human capital distribution is then more compressed among women than among men, as was assumed initially, yielding an overall equilibrium.[4]

An important implication of this mechanism is that, while conventional first-order statistical discrimination has a uniform impact on all members of a group, second-order statistical discrimination has differential impact on different members of the same group. In other words, if two individuals from two groups happen to have identical signals, under first-order statistical discrimination the individual from the advantaged (i.e., higher mean) group will *always* be favored compared to the individual from the disadvantaged (i.e., lower mean) group, regardless of the signal level. On the other hand, under second-order statistical discrimination, a member of the high-variance group could be favored *or* discriminated against, relative to a member of the low-variance group, depending on the value of the signal. Thus, comparing average outcomes across groups may mask potentially severe discrimination, in that discrimination at certain quantiles of the ability or human capital distribution is offset by discrimination in the opposite direction at other quantiles of the distribution.

To give just one example, in 2010 women accounted for 51.5% of employees in management, professional, and related occupations. This ratio suggests that women are not under-represented in these job sectors.[5] However, in the same year women accounted for

therein). According to this explanation, males are evolutionarily conditioned to be more risk seeking than females, resulting in a more variable distribution of traits among males, including intelligence or ability.

[4]At the same time, the mean human capital of the low-variance group may be below, equal to, or above that of the high-variance group. For the described mechanism to work, individuals do not necessarily have to make their investment decisions themselves. The same mechanism would also work if parents made these decisions, as long as the parents' utility depends on their children's labor market success.

[5]If anything, women are slightly over-represented, since they account for only 47.2% of all employed persons. (Source: Bureau of Labor Statistics 2010.)

only 3% of Fortune 500 CEOs.[6] Looking at average employment outcomes for women therefore obscures the severe employment disparities that exist at the far right end of the outcome distribution. Of course, these numbers are not, in and by themselves, evidence of discrimination against women in obtaining top management positions. But they do suggest that, whatever mechanism is causing these employment disparities, the first moment of the outcome distribution does not convey all relevant information contained in this distribution. In the context of statistical discrimination as one possible mechanism, this means that comparing group-wide average labor market outcomes does not allow detection of more subtle forms of discrimination.

The remainder of the paper is organized as follows. In Section 2, we introduce our model. In Section 3, we derive conditions for second-order statistical discrimination and construct one such equilibrium for normally distributed signals. Section 4 concludes with a discussion of our model's relation to the existing literature, its empirical content, and its potential implications for policy.

2 The Model

The population consists of two groups, male and female, denoted $g \in \{m, f\}$. Each group comprises a continuum of measure 1, so the total population has measure 2. Group membership is publicly observable and has no economic significance *ex ante*. That is, all assumptions we make in the model apply equally to both groups.

2.1 Human capital production

Each individual is endowed with an initial ability, which can be either a or b, with $b > a$. The fraction of males and females with ability a is $\lambda \in (0, 1)$. After learning their ability, individuals decide to invest either low effort $\underline{e}$ or high effort $\overline{e}$ in human capital production. Individuals who spend $\underline{e}$ obtain human capital equal to their initial ability. On the other hand, ability a-individuals who invest $\overline{e}$ obtain human capital $A > a$, and ability b-individuals who invest $\overline{e}$ obtain human capital $B > b$. The cost of effort $\underline{e}$ is zero, and the cost of effort $\overline{e}$ is $c > 0$ regardless of ability. An individual's initial ability, effort choice, and final human capital are private information.

The set of possible human capital levels is then $K \equiv \{a, A, b, B\}$. We make two additional assumptions. First, $b > A$: High effort from a low-ability individual is insufficient to overcome a high-ability individual's initial advantage. Second, $B - b > A - a$: Effort has a larger effect on the human capital of high-ability individuals than low-ability individuals.

A strategy for an individual of gender $g \in \{m, f\}$ is a mapping

$$\sigma_g : \{a, b\} \to [0, 1],$$

[6] Source: USA Today 2011.

describing the probability of choosing the high effort level $\overline{e}$ as a function of gender and initial ability. We call the pair $(g, q) \in \{m, f\} \times K$ an individual's type. Given strategy $\sigma = (\sigma_m, \sigma_f)$, the measure of individuals of type (g, q) is given by

$$z_g(q) \;=\; \begin{cases} \lambda(1-\sigma_g(a)) & \text{if } q = a, \\ \lambda\sigma_g(a) & \text{if } q = A, \\ (1-\lambda)(1-\sigma_g(b)) & \text{if } q = b, \\ (1-\lambda)\sigma_g(b) & \text{if } q = B. \end{cases} \tag{1}$$

After individuals have made their effort choices, a publicly observable noisy signal $\theta \in (\underline{\theta}, \overline{\theta}) \subseteq \mathbb{R}$ of an individual's human capital $q \in K$ is generated. The terminology we adopt here is to call θ a "test score." That is, we think of θ as the result of an examination all individuals must take after they have made their effort choices and have acquired their human capital. We make the following assumptions. Conditional on an individual's human capital q, θ has continuous, positive density $f(\theta|q)$ for all $\theta \in (\underline{\theta}, \overline{\theta})$. The corresponding cumulative density is $F(\theta|q)$. To provide information about an individual's human capital, f satisfies the ***monotone likelihood ratio property***:

$$\forall q > q' : \;\; \theta > \theta' \;\Rightarrow\; \frac{f(\theta|q)}{f(\theta|q')} \geq \frac{f(\theta'|q)}{f(\theta'|q')} \;\; (> \text{for some } \theta, \theta'). \tag{MLRP}$$

This property states that, as an individual's human capital increases, high test scores become more likely *relative* to low test scores. Furthermore, f satisfies the following ***separation property***:

$$\lim_{\theta \to \underline{\theta}} \frac{f(\theta|B)}{f(\theta|a)} \;\;<\;\; \frac{\lambda}{1-\lambda}\frac{A-a}{B-b} \;\;<\;\; \lim_{\theta \to \overline{\theta}} \frac{f(\theta|B)}{f(\theta|a)}. \tag{SEP}$$

This property states that test scores in the left (right) tail of the distribution are sufficiently likely (unlikely) to have come from individuals with the lowest human capital a, *relative* to individuals with highest human capital B.

The pair $(g, \theta) \in \{m, f\} \times \mathbb{R}$ will be called the individual's ***public type***. Given the human capital distribution z defined in (1), we can express the density of public type (g, θ) in the population as

$$\widetilde{z}_g(\theta) \;=\; \sum_{q \in K} z_g(q) f(\theta|q).$$

2.2 Job market

There are three different types of jobs in this economy: "Simple jobs" (level 0 jobs), "clerical jobs" (level 1 jobs), and "elite jobs" (level 2 jobs). The measure of available level-i jobs is $\beta_i > 0$ with $\beta_0 + \beta_1 + \beta_2 = 2$ (i.e., every individual can be employed).

We follow Coate and Loury (1993) and assume fixed wages. That is, we assume that individuals attach value V_i to employment in job level i, with $\omega_0 < \omega_1 < \omega_2$.[7]

A precise specification of the demand side structure of the labor market (the number of firms, etc.) is not important for our argument. We only assume that, at the time of hiring, employers observe the individuals' public types, and that employers seek to maximize the expected productivity of their workforce. A labor market outcome is therefore an assignment matching individuals to jobs as a function of their gender and test score. In equilibrium of the labor market, this job-worker assignment must be stable in the following sense: No firm wants to fire one of its current workers and replace him/her with another worker currently employed in a lower ranked job. Since individuals prefer higher ranked jobs over lower ranked ones, any worker who is offered employment in a higher ranked job than his/her current one would accept the higher ranked job offer.

Given these preferences of individuals and employers, a stable assignment must be positive assortative (Becker 1973): Individuals of higher expected productivity are assigned to higher job levels. To formalize this idea, write the expected productivity of an individual of public type (g, θ) as

$$Q_g(\theta) = \frac{\sum_{q \in K} z_g(q) f(\theta|q) \cdot q}{\tilde{z}_g(\theta)} = \frac{\sum_{q \in K} z_g(q) f(\theta|q) \cdot q}{\sum_{q \in K} z_g(q) f(\theta|q)}. \tag{2}$$

Under (MLRP), Q_m and Q_f are strictly increasing in θ (Milgrom 1981). To save on notation, we will set $Q_g(\underline{\theta}) = \lim_{\theta \to \underline{\theta}} Q_g(\theta)$ and $Q_g(\overline{\theta}) = \lim_{\theta \to \overline{\theta}} Q_g(\theta)$. Consider now cutoffs $\hat{\theta}^1_m$ and $\hat{\theta}^1_f$ such that

$$Q_m(\hat{\theta}^1_m) \begin{cases} \geq \\ = \\ \leq \end{cases} Q_f(\hat{\theta}^1_f) \text{ if } \begin{cases} \hat{\theta}^1_f & = \underline{\theta} \\ \hat{\theta}^1_m, \hat{\theta}^1_f & > \underline{\theta} \\ \hat{\theta}^1_m & = \underline{\theta} \end{cases} \quad \text{and} \quad \int_{\underline{\theta}}^{\hat{\theta}^1_m} \tilde{z}_m(\theta) d\theta + \int_{\underline{\theta}}^{\hat{\theta}^1_f} \tilde{z}_f(\theta) d\theta = \beta_0, \tag{3}$$

as well as cutoffs $\hat{\theta}^2_m$ and $\hat{\theta}^2_f$ such that

$$Q_m(\hat{\theta}^2_m) \begin{cases} \geq \\ = \\ \leq \end{cases} Q_f(\hat{\theta}^2_f) \text{ if } \begin{cases} \hat{\theta}^2_f & = \overline{\theta} \\ \hat{\theta}^2_m, \hat{\theta}^2_f & < \overline{\theta} \\ \hat{\theta}^2_m & = \overline{\theta} \end{cases} \quad \text{and} \quad \int_{\hat{\theta}^2_m}^{\overline{\theta}} \tilde{z}_m(\theta) d\theta + \int_{\hat{\theta}^2_f}^{\overline{\theta}} \tilde{z}_f(\theta) d\theta = \beta_2. \tag{4}$$

Condition (3) says that males and females with test scores below $\hat{\theta}^1_m$ and $\hat{\theta}^1_f$, respectively, comprise a measure β_0 of individuals with the lowest expected productivity conditional on their public types. In a stable assignment, these individuals must be assigned to level 0 jobs. Similarly, condition (4) says that males and females with test scores above $\hat{\theta}^2_m$ and $\hat{\theta}^2_f$ comprise a measure β_2 of individuals with the highest expected productivity. In

[7] These values may include wages and monetary benefits, but also non-monetary costs and benefits such as social status or prestige associated with a job, occupational hazards, or the (un)pleasantness of working conditions. The assumption of fixed wages can be replaced with a richer model that features production, coupled with bargaining between workers and firms over the surplus produced in a job.

a stable assignment, these individuals must be assigned to level 2 jobs. All individuals in between these cutoffs fill the remaining measure β_1 of level 1 jobs. Thus, a stable assignment of public types to jobs can be characterized characterized by cutoffs

$$\hat{\theta}^1_m,\, \hat{\theta}^1_f,\, \hat{\theta}^2_m,\, \hat{\theta}^2_f \;\in\; (\underline{\theta}, \overline{\theta})$$

that satisfy conditions (3)–(4) above.[8]

2.3 Equilibrium

An equilibrium of this economy consists of an individual strategy $\sigma = (\sigma_m, \sigma_f)$, an assignment $\hat{\theta} = (\hat{\theta}^1_m, \hat{\theta}^1_f, \hat{\theta}^2_m, \hat{\theta}^2_f)$, and a human capital distribution $z = (z_m, z_f)$ such that three criteria are satisfied. The first is that z is consistent with σ. That is, the equilibrium human capital distribution is generated by (1) from the strategy σ. The second criterion is that the labor market is in equilibrium, that is, the assignment $\hat{\theta}$ is stable. The third criterion is that the individual strategy σ is optimal, given the assignment $\hat{\theta}$. To formalize this third criterion, denote the expected prize for an individual of type (g, q) under assignment $\hat{\theta}$ by

$$W_g(q) \;=\; \omega_0 \int_{\underline{\theta}}^{\hat{\theta}^1_g} f(\theta|q)d\theta \;+\; \omega_1 \int_{\hat{\theta}^1_g}^{\hat{\theta}^2_g} f(\theta|q)d\theta \;+\; \omega_2 \int_{\hat{\theta}^3_g}^{\overline{\theta}} f(\theta|q)d\theta. \tag{5}$$

If the strategy σ is optimal, given $\hat{\theta}$, then the following conditions hold for $g \in \{m, f\}$ and all $s \in [0, 1]$:

$$(1-\sigma_g(a))\, W_g(a) + \sigma_g(a)\big[W_g(A) - c\big] \;\;\geq\;\; (1-s)\, W_g(a) + s\big[W_g(A) - c\big], \tag{6}$$

$$(1-\sigma_g(b))\, W_g(b) + \sigma_g(b)\big[W_g(B) - c\big] \;\;\geq\;\; (1-s)\, W_g(b) + s\big[W_g(B) - c\big]. \tag{7}$$

If these conditions hold, then strategy σ maximizes the individual's expected payoff, given the anticipated assignment $\hat{\theta}$ of workers to jobs.

[8]Note that conditions (3)–(4) imply a productivity threshold above which an individual is eligible for a clerical job, and a second productivity threshold above which an individual is eligible for an elite job. These thresholds (represented by the two horizontal lines in Figure 1 in Section 3) are endogenous and adjust so that the correct mass of workers β_i is employed in each sector. This causal direction can be reversed. That is, one could consider an alternative specification of the labor market where expected productivity must be above an exogenous threshold $\overline{Q}_i$ in order to be eligible for a level-i job, and where the mass of workers employed in each sector adjusts endogenously. These specifications are outcome-equivalent: For each pair of exogenous capacity constraints (β_1, β_2) in the original model, there exists a pair of exogneous productivity thresholds $(\overline{Q}_1, \overline{Q}_2)$ that results in the same signal cutoffs $\hat{\theta}^1_m, \hat{\theta}^1_f, \hat{\theta}^2_m, \hat{\theta}^2_f$ (and vice versa). The bijection between (β_1, β_2) and $(\overline{Q}_1, \overline{Q}_2)$ is given by (3)–(4), with $\overline{Q}_1 \equiv \min\{Q_m(\hat{\theta}^1_m), Q_f(\hat{\theta}^1_f)\}$ and $\overline{Q}_2 \equiv \min\{Q_m(\hat{\theta}^2_m), Q_f(\hat{\theta}^2_f)\}$.

An equilibrium $(\sigma, \hat{\theta}, z)$ is ***non-discriminatory*** if $\hat{\theta}_m^i = \hat{\theta}_f^i$ for $i = 1, 2$. Otherwise, it is called ***discriminatory***. In a discriminatory equilibrium, a different minimum test score is required for males, compared to females, in order to qualify for clerical and/or elite jobs. That is $\hat{\theta}_m^1 \neq \hat{\theta}_f^1$ or $\hat{\theta}_m^2 \neq \hat{\theta}_f^2$ or both.

Before characterizing equilibrium in the next section, we establish its existence. (Proofs of all results are in the Appendix.)

Proposition 1. *An equilibrium exists. In particular, a non-discriminatory equilibrium exists.*

3 Discriminatory Equilibrium

In this section we explore the possibility of discriminatory equilibrium. We focus on a particular equilibrium, which features the following pure strategy for individuals:

$$\sigma_m(a) = \sigma_f(b) = 0, \;\; \sigma_m(b) = \sigma_f(a) = 1. \tag{8}$$

That is, low-ability males and high-ability females exert low effort, while high-ability males and low-ability females exert high effort. The human capital distribution in the population is then given by

$$z_m(a) = z_f(A) = \lambda, \;\; z_m(B) = z_f(b) = 1 - \lambda. \tag{9}$$

Note that, in our candidate equilibrium, the female human capital distribution is more compressed than the male human capital distribution. As we will show below, (9) implies that the equilibrium job-worker assignment $\hat{\theta}$ will be such that

$$\hat{\theta}_f^1 \; < \; \hat{\theta}_m^1 \; < \; \hat{\theta}_m^2 \; < \; \hat{\theta}_f^2. \tag{10}$$

Female workers are hence disadvantaged relative to males when it comes to obtaining an elite job, in the sense that the test score of a female worker must meet a higher threshold requirement in order to get an elite job. The opposite holds for the cutoff score needed to obtain a clerical job, where male workers are disadvantaged vis-á-vis female workers.

We will show that, under certain conditions, anticipation of the labor market assignment (10) will discourage women of high ability as well as men of low ability from investing into their human capital, yielding the strategy (8) and thus the human capital distribution (9). In this situation, more males than females will be employed in the elite sector and in simple jobs, and more females than males will be employed in clerical jobs. We call this outcome ***second-order statistical discrimination***. It should be clear that, by relabeling, one can construct another equilibrium in which the male and female roles are reversed.

3.1 Conditions for discriminatory equilibria

Using the human capital distribution in (9), the Bayesian posterior probability that a male worker with test score θ has human capital a or B is

$$Pr[a|m,\theta] = \frac{\lambda f(\theta|a)}{\lambda f(\theta|a) + (1-\lambda) f(\theta|B)}, \quad Pr[B|m,\theta] = \frac{(1-\lambda) f(\theta|a)}{\lambda f(\theta|a) + (1-\lambda) f(\theta|B)}.$$

(The human capital levels A and b must receive a zero probability for male workers.) The expected productivity of a male worker with test score θ is therefore given by

$$Q_m(\theta) = \frac{\lambda f(\theta|a) a + (1-\lambda) f(\theta|B) B}{\lambda f(\theta|a) + (1-\lambda) f(\theta|B)}.$$

The expected productivity of this worker can be similarly written as

$$Q_f(\theta) = \frac{\lambda f(\theta|A) A + (1-\lambda) f(\theta|b) b}{\lambda f(\theta|A) + (1-\lambda) f(\theta|b)}.$$

Because f is continuous in θ, the expectations Q_m and Q_f are continuous on $(\underline{\theta}, \overline{\theta})$. Furthermore:

Lemma 2. *Q_m and Q_f are increasing in θ and satisfy*

$$\lim_{\theta \to \underline{\theta}} Q_m(\theta) \;>\; \lim_{\theta \to \underline{\theta}} Q_f(\theta) \;>\; \lim_{\theta \to \overline{\theta}} Q_f(\theta) \;>\; \lim_{\theta \to \overline{\theta}} Q_m(\theta).$$

Because Q_m and Q_f are continuous in θ, Lemma 2 implies that they must intersect at least once. Let θ^* be the left-most intersection of Q_m and Q_f, and let θ^{**} be the right-most intersection. To construct the equilibrium assignment $\hat{\theta}$, we consider the following scenario:

$$\int_{\underline{\theta}}^{\theta^*} \big(\widetilde{z}_m(\theta) + \widetilde{z}_f(\theta)\big) d\theta \;>\; \beta_0, \tag{11}$$

$$\int_{\theta^{**}}^{\overline{\theta}} \big(\widetilde{z}_m(\theta) + \widetilde{z}_f(\theta)\big) d\theta \;>\; \beta_2. \tag{12}$$

Condition (11) states that the fraction of males and females with test scores below θ^* exceeds the capacity β_0. Similarly, condition (12) states that the fraction of males and females with test scores above θ^{**} exceeds the capacity β_2. Under these conditions, both males and females will be employed in the clerical sector. Some males, and possibly some females, will be employed in the elite sector as well as in the simple sector.

Cutoff thresholds $(\hat{\theta}^1_m, \hat{\theta}^1_f,\, \hat{\theta}^2_m, \hat{\theta}^2_m)$ to place exactly β_i individuals into each level i job can now be computed by applying conditions (3)–(4). Because $Q_m(\theta) < Q_f(\theta)$ for $\theta < \theta^*$ by Lemma 2, (3) together with (11) implies $\hat{\theta}^1_f < \hat{\theta}^1_m < \theta^*$. Similarly, because $Q_m(\theta) > Q_f(\theta)$ for $\theta < \theta^{**}$, (4) together with (12) implies $\theta^{**} < \hat{\theta}^2_m < \hat{\theta}^2_f$. Therefore, the assignment satisfies the inequalities (10). Figure 1 illustrates its construction graphically. The heights of the two horizontal lines are defined implicitly by (3) and (4).

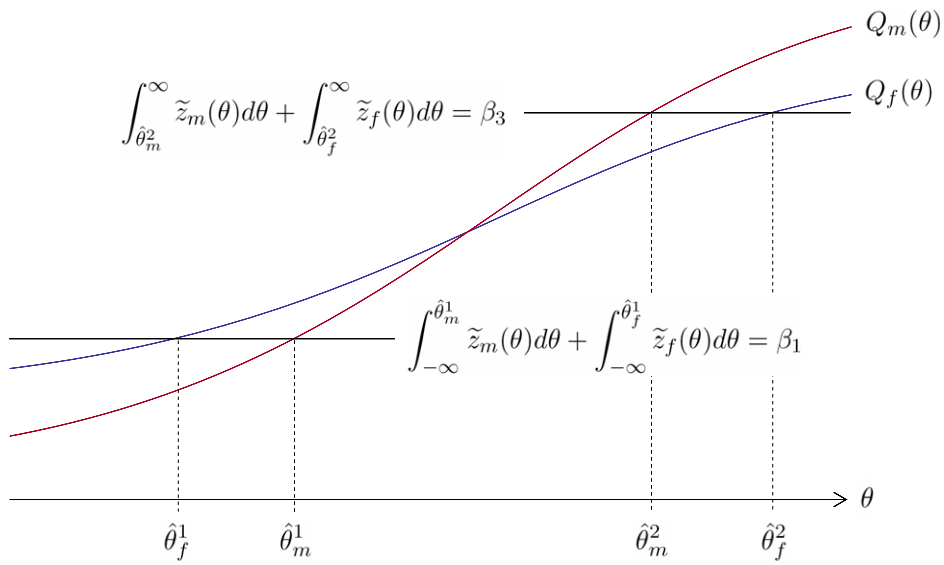

Figure 1: Assignment in discriminatory equilibrium

Finally, we must check whether the equilibrium effort strategy posited in (8) is indeed optimal for each individual, given the assignment constructed above. Using (5)–(7), low effort is optimal for male workers with innate ability a and high effort is optimal for male workers with innate ability b if and only if

$$\begin{aligned}(\omega_1-\omega_0)\big[F(\hat\theta_m^1|a)-F(\hat\theta_m^1|A)\big] \;+\; (\omega_2-\omega_1)\big[F(\hat\theta_m^2|a)-F(\hat\theta_m^2|A)\big]\\ \le\; c \;\le\; (\omega_1-\omega_0)\big[F(\hat\theta_m^1|b)-F(\hat\theta_m^1|B)\big] \;+\; (\omega_2-\omega_1)\big[F(\hat\theta_m^2|b)-F(\hat\theta_m^2|B)\big].\end{aligned} \tag{13}$$

The reverse inequalities hold for female workers:

$$\begin{aligned}(\omega_1-\omega_0)\big[F(\hat\theta_f^1|a)-F(\hat\theta_f^1|A)\big] \;+\; (\omega_2-\omega_1)\big[F(\hat\theta_f^2|a)-F(\hat\theta_f^2|A)\big]\\ \ge\; c \;\ge\; (\omega_1-\omega_0)\big[F(\hat\theta_f^1|b)-F(\hat\theta_f^1|B)\big] \;+\; (\omega_2-\omega_1)\big[F(\hat\theta_f^2|b)-F(\hat\theta_f^2|B)\big].\end{aligned} \tag{14}$$

A discriminatory equilibrium therefore exists if (11)–(12) as well as (13)–(14) are satisfied, where the assignment $\hat\theta = (\hat\theta_m^1, \hat\theta_f^1, \hat\theta_m^2, \hat\theta_f^2)$ is derived from (3)–(4).

3.2 Gaussian test score distribution

In Section 3.1 we derived sufficient conditions for a discriminatory equilibrium in which the female human capital distribution had a lower variance than the corresponding male distribution. To say more, we now assume that for an individual with human capital

$q \in K$, the score θ is distributed normally on $(-\infty, \infty)$ with mean q and variance ν^2. That is, an individual's test score is the sum of his or her human capital and Gaussian noise. The conditional density $f(\theta|q)$ is given by

$$f(\theta|q) \;=\; \frac{1}{\sqrt{2\pi}\nu} e^{-\frac{1}{2\nu^2}(\theta - q)^2}. \tag{15}$$

Note that the Gaussian distribution satisfies both the monotone likelihood ratio property (MLRP) and the separation property (SEP) we assumed throughout.

We now show, by means of an example, that equilibria with second-order statistical discrimination generically exist when test scores are normally distributed.

Example 1. Suppose that test scores are normally distributed with variance $\nu^2 = 1/16$, and let the other parameters of the model take on the following values:

$$a = 0, \;\; A = 0.3, \;\; b = 0.6, \;\; B = 1, \;\; \beta_0 = 1, \;\; \beta_1 = 0.75, \;\; \beta_2 = 0.25,$$
$$\omega_0 = 0, \;\; \omega_1 = 1, \;\; \omega_2 = 1.5, \;\; \lambda = 0.8, \;\; c = 0.4.$$

With these parameter values, Q_m and Q_f cross once at 0.559. The inequalities (11)–(12) are satisfied:

$$\int_{-\infty}^{0.559} \big(\widetilde{z}_m(\theta) + \widetilde{z}_f(\theta)\big) d\theta = 1.564 \, > \, \beta_0, \quad \int_{0.559}^{\infty} \big(\widetilde{z}_m(\theta) + \widetilde{z}_f(\theta)\big) d\theta = 0.436 \, > \, \beta_2.$$

Constructing the cutoff scores for a stable assignment as in (3)–(4), we get

$$\hat{\theta}^1_f = 0.131, \; \hat{\theta}^1_m = 0.538, \; \hat{\theta}^2_m = 0.582, \; \hat{\theta}^2_f = 0.824.$$

The individual optimality condition for males is now $0.215 < c < 0.582$, and for females it is $0.459 > c > 0.317$, which implies that individuals in fact follow strategy (8). Since conditions (11)–(12) and (13)–(14) hold as strict inequalities, an equilibrium with second-order statistical discrimination exists for a non-empty open set of parameters. The proportion of males and females employed in each of the three job sectors is as follows:

	Males	Females	Total
Elite	0.199	0.051	0.25
Clerical	0.008	0.742	0.75
Simple	0.794	0.206	1

Furthermore, the average equilibrium human capital (and test score) for male workers this example is 0.2; while for female workers it is 0.36. It is hence possible that one group is more able on average, while another group has a higher maximum ability and will occupy a larger share of elite jobs.

With perfect information about an individual's skills, of course, statistical discrimination (of any kind) cannot arise. One may therefore suspect that, if the signals about one's qualifications become more precise, the possibility of second-order statistical discrimination ceases to exist. For normally distributed test scores, however, quite the opposite can be the case:

Proposition 3. *Assume the conditional test score distribution is normal with mean* $q \in K$ *and variance* $\nu^2 > 0$. *Suppose further that* $\beta_0 \in (\lambda, 2\lambda)$, $\beta_2 \in (0, 1-\lambda)$, *and*

$$0 < c < \min\left\{(\omega_1 - \omega_0)\bigl(2 - \beta_0/\lambda\bigr), (\omega_2 - \omega_1)\beta_2/(1-\lambda)\right\}.$$

Then there exists $\overline{\nu}^2 > 0$ *such that second-order statistical discrimination is an equilibrium outcome for all* $\nu^2 \in \bigl(0, \overline{\nu}^2\bigr)$. *In any such equilibrium, only male workers will be employed in elite jobs.*

Since access to better information improves the matching of workers to jobs, firms seem to have an incentive to invest in the accuracy of signals they obtain about the qualifications of potential hires. As long as some arbitrarily small uncertainty about workers' qualifications remains, however, Proposition 3 implies that increasing the accuracy of signals alone does not guarantee the elimination of second-order statistical discrimination in the labor market. In particular, there is no guarantee that the "glass ceiling" disappears as signals become increasingly accurate.

4 Discussion

We extended the basic framework of self-fulfilling statistical discrimination from the first moment of the human capital distribution to the second. We now conclude with a discussion of how our theory relates to the existing literature on statistical discrimination, how it can be distinguished empirically from first-order theories of statistical discrimination, and what it might imply for anti-discrimination policies.

4.1 Relation to the theoretical literature

Early models of statistical discrimination assume that the conditional signal about an individual's qualification is more precise for one group than for another (Phelps 1972; Aigner and Cain 1977). In this case, employers place relatively more weight on group averages, and relatively less weight on individual signals, when forming expectations about the qualifications of applicants from the group with less precise signals (women, say) compared to applicants with more precise signals (men, say). Three related phenomena arise in this setup. First, the human capital distribution among men has a higher variance than among women. Second, the same is true for the labor market outcomes of men and women. Third, discrimination has a non-uniform impact on workers at different locations on the skill or signal spectrum: Women with a high signal earn less than men with the same signal, while men with a low signal earn less than women with the same signal.[9]

Our discriminatory equilibrium is characterized by the same three features. However, these features emerge as the result of a very different mechanism, relying on a different set of assumptions. Most importantly, in our framework the two groups are assumed to be symmetric *ex ante*: Except for the gender label, there are no differences between men and women (in particular, there are no differences in signaling technology across groups). Instead, to generate its result our model relies on the interplay of two further assumptions: Within-group heterogeneity in endowed ability, and a three-task assignment specification of the labor market. All three assumptions have previously been explored in the literature on statistical discrimination, but not in conjunction.

First, consider endowment heterogeneity. It is well-known that the assumption of within-group heterogeneity in ability can generate second-moment differences in the human capital distribution across groups. For example, this is the case already in Phelps (1972), where it also leads to second-moment differences in competitive wages and non-uniform impacts of discrimination at different ends of the signal spectrum—at the expense of having to assume exogenous differences in signaling technologies. Note, however, that the main purpose of unobserved endowment heterogeneity in Phelps (1972) is to create a meaningful signal extraction problem for employers. Coate and Loury (1993) show that unobserved human capital investments can play the same role, so that a similar signal extraction problem can be constructed with homogeneous endowments. Together with the assumption of a two-task/fixed-wage assignment model, this allows Coate and Loury (1993) to do away with any exogenous inter-group differences, and instead replace these with a self-fulfilling expectations theory of endogenous inter-group differences. On the other hand, the analytical focus has now shifted purely to first-order statistical discrimination. Moro (2003), finally, extends Coate and Loury's (1993) model to again allow for heterogenous endowments, thus obtaining second-moment differences in the human capital and the outcome distribution, without necessarily having to assume differences in signaling technologies. However, unlike in Phelps' (1972) original model, this does not translate into a differential impact of discrimination across the endowment spectrum.

The reason is that Moro (2003) maintains the two-task assignment model of the labor market introduced by Coate and Loury (1993). Having two tasks means that individuals strive for only one type of "prize" (i.e., being assigned to the better of the two tasks). This closes the channel through which heterogenous endowments can translate into second-moment differences in outcomes. In our model, on the other hand, a third task allows for a richer form of competition: The middle job essentially becomes the prize over which individuals with low endowments compete, while the top job becomes the prize over which individuals with high endowments compete. Second-moment differences in the

[9]If a human capital investment stage is added, the expected return of investments is lower for women than for men, which also depresses the *average* human capital among women, compared to men (see Lundberg and Startz 1983).

human capital distribution across genders can now manifest themselves in second-order statistical discrimination in the labor market.[10]

Finally, the assumption of three job levels is also used in Bjerk (2008). In this career ladder model, a glass-ceiling effect arises for women even though there is no discrimination in promotion to top-level jobs. Instead, the glass ceiling is the result of statistical discrimination earlier in a worker's career, which delays promotion of female employees from entry-level jobs to mid-level jobs. This statistical discrimination, however, is caused by assumed exogenous differences in the precision of skill signals across groups. In other words, in this model groups are not *ex ante* identical.

4.2 Empirical implications

As we have shown theoretically, the presence of discrimination may remain undetected when comparing average outcomes across groups. As is often the case, the first moment of a distribution does not convey all relevant information, so that a comparison of group-wide average labor market outcomes can mask more subtle forms of discrimination. A growing empirical literature is concerned with this issue.

Jenkins (1994) points out the deficiencies of using only the mean gender wage gap as a measure of discrimination, and proposes new measurements that capture discrimination across the whole distribution. Using 1980 UK earnings data, he finds evidence that discrimination against women strictly increases when moving from the lower quartile to the median to the upper quartile. Albrecht *et al.* (2003) find a similar pattern in 1998 Swedish data. Using 1995 and 1999 Spanish data, del Rio *et al.* (2006) and de la Rica *et al.* (2008) show that for college-educated women, the gender wage gap is wider at the upper tail of the wage distribution, while it is wider at the lower tail for less educated women. Lastly, Arulampalam *et al.* (2007) document similar patterns for these and other European countries for the period 1995–2001.

Of course, these patterns can have many explanations. Assuming (statistical) discrimination as one underlying cause, is it possible to empirically distinguish the presence of second-order discrimination from discrimination of the first order? The answer is a qualified yes. The fundamental difference between first-order and second-order statistical discrimination is that the former has a uniform impact on all members of the disadvantaged group (i.e., all members of a group are hurt by discrimination), while the latter can have qualitatively different impacts on different members of the same groups (i.e., some members of a group are hurt by discrimination, but others benefit). Thus, a good indicator of second-order discrimination would be a reversal of the inter-group outcome gap. For example, a positive gender wage gap at the bottom quantiles of the earnings distribution in a labor market, together with a negative wage gap at the top quantiles, would strongly suggest the presence of second-order discrimination.

[10] It should be clear that this logic continues to apply for any number of job levels or tasks, as long as it is larger than two.

While the aforementioned empirical studies document changes in the female wage gap across quantiles, they do not find evidence of a gap reversal. This does not mean that second-order discrimination does not play a role, as it could coexist with first-order discrimination, complicating the task of cleanly separating the consequences of each kind of discrimination in the data. For example, if statistical discrimination results in a larger variance of male earnings compared to female earnings, and at the same time systematically shifts the female earnings distribution to the left, the latter effect may override any observable gap reversals the former effect generates.

Reversal effects have been documented, however, in other contexts. For the case of racial groups, Bjerk (2007) uses Armed Forces Qualification Test (AFQT) scores as a measure for workers' pre-market skills, and finds significant differences in skill-sorting thresholds across racial groups: While black workers are less likely than white workers to be employed in white-collar jobs, conditional on AFQT scores blacks are *more* likely to work in white-collar jobs. In particular, the thresholds that sort black workers into white-collar job sectors are significantly lower than those for white workers. If employers believe that the black group has a higher skill variance than the white group, then, conditional on a high AFQT score, a black worker has a higher expected productivity than a white worker with the same score, and is hence favored by potential employers in white-collar job sectors. Thus, individual black workers at the high end of the ability distribution can benefit from statistical discrimination, despite the fact that blacks have a lower average skill level than whites.

Finally, Chiswick *et al.* (2008) analyze the immigrant wage gap in the U.S. and Australia and find that, in both cases, immigrants from non-English speaking countries have an advantage over native-born workers at the low end of the earnings distribution, but a disadvantage at the upper end, everything else equal. This pattern can be reconciled with second-order statistical discrimination but is inconsistent with fist-order statistical discrimination. However, this study does not directly answer whether the variance difference between immigrants and native-born workers are due to self-fulfilling expectations or other reasons, such as self-selection in the immigration decisions.

4.3 Policy aspects

Like every model of statistical discrimination, ours relies on an informational friction which prevents employers from learning an individual's skills perfectly. This friction induces a two-fold inefficiency: First, it discourages some high-ability individuals from investing, while some low-ability individuals invest. Because the cost of investing is assumed to be constant but the returns are larger for high-ability individuals, this is a misallocation of resources. Second, taken the investment decisions as given, the final allocation of workers to jobs may not be optimal. It is questionable whether markets can correct these efficiencies by themselves. As shown in Proposition 3, even when signals become increasingly precise it is possible for second-order statistical discrimination to remain an equilibrium. Thus, some interventions are needed to correct the inefficiencies resulting from it. It is beyond the scope of this work to examine the many conceivable policy measures that could be used, but our results allow us to at least speculate on some aspects of policy.

First, the differential impact of second-order statistical discrimination on individuals within the same group suggests that successful policies may have to be targeted at certain jobs, individuals, or firms only. For example, a broad-stroke mandate that firms employ an equal number of men and women would have no effect on firms that offer all three tiers of jobs. On the other hand, a mandate that an equal number of men and women be employed in elite jobs would have an equalizing effect on male and female employment in this sector (and might then result in altered investment incentives for high-ability females). Second, policies targeting the elite job sector may have spill-over effects to lower job sectors. An equal-employment mandate on the elite job level will reallocate some women from the clerical sector to the top sector, and some men from the top sector the clerical sector. These changes could, in turn, affect the assignment of workers to the simple sector. Thus, even a highly targeted policy intervention can potentially affect the labor market outcomes of all individuals, and hence the investment incentives of all individuals.

Appendix

Proof of Proposition 1

Without loss of generality we can take the support of the test score distribution f to be a subset of $[0,1]$, if necessary by constructing a new test score $t \in [0,1]$ and setting $t \equiv e^{\theta}/(1+e^{\theta})$. An equilibrium is then a point $(\sigma, \hat{\theta}, z) \in S \equiv [0,1]^4 \times [0,1]^4 \times \Delta^2$ satisfying the equilibrium conditions in Section 2.3, where

$$\Delta \equiv \Big\{ \big(z(a), z(A), z(b), z(B)\big) \in [0,1]^4 \, : \, z(a)+z(A)+z(b)+z(B)=1 \Big\}$$

is the set of human capital distributions within each group. Define three correspondences

$$T^{\sigma} : [0,1]^4 \rightarrow\rightarrow [0,1]^4, \quad T^{\hat{\theta}} : \Delta^2 \rightarrow\rightarrow [0,1]^4, \quad T^{z} : [0,1]^4 \rightarrow\rightarrow \Delta^2$$

as follows:

$$\begin{array}{rcll} T^{\sigma}(\hat{\theta}) & = & \{ \ \sigma' \in [0,1]^4 \ : \ s' \text{ is optimal given } \hat{\theta} & \}, \\ T^{\hat{\theta}}(z) & = & \{ \ \hat{\theta}' \in [0,1]^4 \ : \ \hat{\theta}' \text{ is stable given } z & \}, \\ T^{z}(\sigma) & = & \{ \ z' \in \Delta^2_K \ : \ z' \text{ is consistent with } \sigma & \}. \end{array}$$

The definitions of optimality, stability, and consistence (given in Section 2.2 and 2.3) imply that T^{σ}, $T^{\hat{\theta}}$, T^{z} are all upper-hemicontinuous. Define a new correspondence $T : S \rightarrow\rightarrow S$ by setting $T(\sigma, \hat{\theta}, z) \equiv T^{\sigma}(\hat{\theta}) \times T^{\hat{\theta}}(z) \times T^{z}(\sigma)$. T is upper-hemicontinuous on a compact and convex set $S \subset \mathbb{R}^{16}$. By Kakutani's Fixed Point Theorem there exists a point $(\sigma, \hat{\theta}, z) \in T(\sigma, \hat{\theta}, z)$, satisfying our equilibrium conditions.

To show that a non-discriminatory equilibrium exists, let $\hat{S} \subset S$ be defined as

$$\hat{S} \equiv \left\{ (\sigma, \hat{\theta}, z) \in S \,:\, \sigma_m = \sigma_f,\, \hat{\theta}^1_m = \hat{\theta}^1_f,\, \hat{\theta}^2_m = \hat{\theta}^2_f,\, z_m = z_f \right\}.$$

Observe that $\hat{S}$ is a compact and convex subset of $\mathbb{R}^{16}$. Furthermore, $T(\sigma, \hat{\theta}, z) \cap \hat{S} \neq \emptyset$ for all $(\sigma, \hat{\theta}, z) \in \hat{S}$. To see this, make the following observations: (i) If $\hat{\theta}^1_m = \hat{\theta}^1_f$ and $\hat{\theta}^2_m = \hat{\theta}^2_f$, an effort decision is optimal for males if and only if the same decision is optimal for females; thus there exists an optimal strategy $s = (s_m, s_f)$ with $s_m = s_f$. (ii) If $s_m = s_f$, the human capital distribution among males is the same as among females: $z_m = z_f$. (iii) If $z_m = z_f$, then $Q_m = Q_f$ and a stable assignment in the labor market will be such that $\hat{\theta}^1_m = \hat{\theta}^1_f$ and $\hat{\theta}^2_m = \hat{\theta}^2_f$. Thus, T can be restricted to an upper-hemicontinuous correspondence from $\hat{S}$ into $\hat{S}$. By Kakutani's Fixed Point Theorem, a fixed point in $\hat{S}$ exists, which then satisfies the definition of a non-discriminatory equilibrium. □

Proof of Lemma 2

Define

$$\gamma_m(\theta) \equiv \left[1 + \frac{1-\lambda}{\lambda} \frac{f(\theta|B)}{f(\theta|a)}\right]^{-1}, \quad \gamma_f(\theta) \equiv \left[1 + \frac{1-\lambda}{\lambda} \frac{f(\theta|b)}{f(\theta|A)}\right]^{-1}.$$

We first show that, under (MLRP), $\gamma_m(\theta)$ and $\gamma_f(\theta)$ satisfy the following properties:

(i) $\gamma_m(\theta)$ and $\gamma_f(\theta)$ are weakly decreasing in θ,

(ii) $\lim_{\theta \to \underline{\theta}} \gamma_m(\theta) > \lim_{\theta \to \overline{\theta}} \gamma_m(\theta)$ and $\lim_{\theta \to \underline{\theta}} \gamma_f(\theta) > \lim_{\theta \to \overline{\theta}} \gamma_f(\theta)$,

(ii) $\lim_{\theta \to \underline{\theta}} \gamma_m(\theta) > \lim_{\theta \to \underline{\theta}} \gamma_f(\theta)$ and $\lim_{\theta \to \overline{\theta}} \gamma_m(\theta) > \lim_{\theta \to \overline{\theta}} \gamma_f(\theta)$.

To show (i), note that (MLRP) implies $f(\theta|B)/f(\theta|a)$ and $f(\theta|b)/f(\theta|A)$ are increasing in θ. Thus, $\gamma_m(\theta)$ and $\gamma_f(\theta)$ are decreasing in θ. To show (ii), suppose $\lim_{\theta \to \underline{\theta}} \gamma_m(\theta) \leq \lim_{\theta \to \overline{\theta}} \gamma_m(\theta)$. By claim (i), this implies $\gamma_m(\theta)$ is a constant for all θ, which in turn implies that $f(\theta|B)/f(\theta|a)$ is independent of θ. This is not possible, due to the stricty inequality part of (MLRP). To show (iii), express $f(\theta|B)/f(\theta|a)$ as

$$\frac{f(\theta|B)}{f(\theta|a)} = \frac{f(\theta|b)}{f(\theta|A)} \cdot \frac{f(\theta|A)}{f(\theta|a)} \frac{f(\theta|B)}{f(\theta|b)}$$

and observe that, as θ decreases, $f(\theta|A)/f(\theta|a)$ and $f(\theta|B)/f(\theta|b)$ are decreasing due to (MLRP). Furthermore, both terms must fall below one for θ small enough. To see why, suppose to the contrary that $f(\theta|A)/f(\theta|a) \geq 1$ as $\theta \to \underline{\theta}$. This implies $f(\theta|A) \geq f(\theta|a)$ for all θ due to (MLRP). Furthermore, $f(\theta|A) < f(\theta|a)$ on an open set of values for θ, due to the strict inequality part of (MLRP) and continuity of f. But then $1 = \int f(\theta|A) d\theta > \int f(\theta|a) d\theta = 1$, a contradiction. The same argument applies to $f(\theta|B)/f(\theta|b)$. It follows that $f(\theta|B)/f(\theta|a) < f(\theta|b)/f(\theta|A)$ as $\theta \to \underline{\theta}$, and therefore $\lim_{\theta \to \underline{\theta}} \gamma_m(\theta) > \lim_{\theta \to \underline{\theta}} \gamma_f(\theta)$. The second inequality can be shown in the same fashion.

Now express Q_m and Q_f as follows:

$$Q_m(\theta) = \gamma_m(\theta)a \;+\; [1-\gamma_m(\theta)]\,B, \qquad Q_f(\theta) = \gamma_f(\theta)A \;+\; [1-\gamma_f(\theta)]\,b. \tag{16}$$

Since $a < B$ and $A < b$, claim (i) implies Q_m and Q_f are increasing in θ. Next, we show that

$$\lim_{\theta\to\underline{\theta}} Q_m(\theta) \;<\; \lim_{\theta\to\underline{\theta}} Q_f(\theta) \;<\; \lim_{\theta\to\overline{\theta}} Q_f(\theta) \;<\; \lim_{\theta\to\overline{\theta}} Q_m(\theta).$$

The middle inequality follows from claim (ii). To show the left inequality, note that (SEP) implies

$$1 + \frac{1-\lambda}{\lambda} \lim_{\theta\to\underline{\theta}} \frac{f(\theta|B)}{f(\theta|a)} \;<\; 1 + \frac{A-a}{B-b}$$

and thus $\lim_{\theta\to\underline{\theta}} \gamma_m(\theta)\cdot(B-b+A-a) > B-b$. By claim (iii) and $b > A$ it follows that $\lim_{\theta\to\underline{\theta}} \gamma_m(\theta)\cdot(B-a) - \lim_{\theta\to\underline{\theta}} \gamma_f(\theta)\cdot(b-A) > B-b$. Rearranging, we get

$$\lim_{\theta\to\underline{\theta}} \gamma_m(\theta)\cdot a + \left[1 - \lim_{\theta\to\underline{\theta}} \gamma_m(\theta)\right]\cdot B \;<\; \lim_{\theta\to\underline{\theta}} \gamma_f(\theta)\cdot A + \left[1 - \lim_{\theta\to\underline{\theta}} \gamma_f(\theta)\right]\cdot b$$

or $\lim_{\theta\to\underline{\theta}} Q_m(\theta) < \lim_{\theta\to\underline{\theta}} Q_f(\theta)$, as desired. The inequality $\lim_{\theta\to\overline{\theta}} Q_f(\theta) < \lim_{\theta\to\overline{\theta}} Q_m(\theta)$ can be shown in a similar manner. □

Proof of Proposition 3

Express $Q_m(\theta)$ and $Q_f(\theta)$ as in (16). Using the normal distribution for θ, the weights $\gamma_m(t)$ and $\gamma_f(\theta)$ can be written as

$$\gamma_m(\theta) \;=\; \left[1 + \frac{1-\lambda}{\lambda}\exp\Bigl(-\frac{1}{2\nu^2}(B-a)(B+a-2\theta)\Bigr)\right]^{-1},$$
$$\gamma_f(\theta) \;=\; \left[1 + \frac{1-\lambda}{\lambda}\exp\Bigl(-\frac{1}{2\nu^2}(b-A)(b+A-2\theta)\Bigr)\right]^{-1}.$$

Note that that $\gamma_m((a+B)/2) = \gamma_f((A+b)/2) = \lambda$. Observe further that for every $\delta > 0$ one can find $\varepsilon > 0$ such that $\nu^2 < \varepsilon$ implies $\gamma_m(\theta) > 1-\delta$ for all $\theta < (B+a)/2 - \delta$, as well as $\gamma_m(\theta) < \delta$ for all $\theta > (B+a)/2 + \delta$. Similarly, for every $\delta > 0$ one can find $\varepsilon > 0$ such that $\nu^2 < \varepsilon$ implies $\gamma_f(\theta) > 1-\delta$ for all $\theta < (b+A)/2 - \delta$, as well as $\gamma_f(\theta) < \delta$ for all $\theta > (b+A)/2 + \delta$. As $\nu^2 \to 0$, therefore, the expectations Q_m and Q_f converge pointwise to

$$\lim_{\nu^2\to 0} Q_m(\theta) \;=\; \begin{cases} a & \text{if } \theta < (a+B)/2, \\ \lambda a + (1-\lambda)B & \text{if } \theta = (a+B)/2, \\ B & \text{if } \theta > (a+B)/2 \end{cases}$$

and

$$\lim_{\nu^2\to 0} Q_f(\theta) \;=\; \begin{cases} A & \text{if } \theta < (A+b)/2, \\ \lambda A + (1-\lambda)b & \text{if } \theta = (A+b)/2, \\ b & \text{if } \theta > (A+b)/2. \end{cases}$$

Note that, for $\nu^2 \approx 0$, Q_m and Q_f intersect once at $\theta^* \approx (a+B)/2$. Note also that $B-b > A-a$ implies $(a+B)/2 > (A+b)/2$, and that $b > A$ implies $(A+b)/2 > A$.

Next, we construct the job-worker assignment $\hat{\theta}$ and examine its properties for small ν^2. Throughout, if x and y are two variables, the statement "$x \approx y$ for $\nu^2 \approx 0$" means "$\forall\delta > 0\ \exists\varepsilon > 0$ s.t. $\nu^2 < \varepsilon \Rightarrow |x-y| < \delta$." Note that the distribution of test scores, $(\widetilde{z}_m, \widetilde{z}_f)$, is concentrated around the human capital levels a, A, b, B, and increasingly so as $\nu^2 \to 0$. For $\nu^2 \approx 0$, therefore, we have

$$\int_{-\infty}^{\theta^*} \big(\widetilde{z}_m(\theta) + \widetilde{z}_f(\theta)\big) d\theta \;\approx\; 2\lambda, \quad \int_{\theta^*}^{\infty} \big(\widetilde{z}_m(\theta) + \widetilde{z}_f(\theta)\big) d\theta \;\approx\; 2(1-\lambda).$$

Given $\beta_0 \in (\lambda, 2\lambda)$ and $\beta_2 \in (0, 1-\lambda)$, conditions (11)–(12) hold for sufficiently small ν^2. Furthermore, the condition $\beta_0 \in (\lambda, 2\lambda)$ means that there are more simple jobs than males with human capital a, but fewer than males with human capital a and females with human capital A combined. For $\nu^2 \approx 0$, this implies that approximately λ males and approximately $\beta_0 - \lambda < \lambda$ of females will be employed in the simple sector. The cutoff scores that separate simple jobs from clerical jobs therefore satisfy $\hat{\theta}_f^1 \approx A < (A+b)/2 < (a+B)/2 \approx \hat{\theta}_m^1$. On the other hand, the condition $\beta_2 \in (0, 1-\lambda)$ means that there are fewer elite jobs than males with human capital B. For $\nu^2 \approx 0$, this implies that exactly β_2 males and no females will be employed in the elite sector. The cutoff scores that separate clerical jobs from elite jobs therefore satisfy $(B+a)/2 < B \approx \hat{\theta}_m^2 < \hat{\theta}_f^2 = \infty$. For small enough ν^2, the job-worker assignment therefore satisfies $\hat{\theta}_f^1 < \hat{\theta}_m^1 < \hat{\theta}_m^2 < \hat{\theta}_f^2$, as required in a second-order discriminatory equilibrium.

Finally, we examine the individual incentives to invest effort, under the assignment constructed above and again assuming that $\nu^2 \approx 0$. First, consider males of initial ability a. For $\nu^2 \approx 0$, we have $F(\hat{\theta}_m^1|a) \approx F(\hat{\theta}_m^1|A) \approx F(\hat{\theta}_m^2|a) \approx F(\hat{\theta}_m^2|A) \approx 1$. The expression on the left side of (13) therefore becomes

$$(\omega_1 - \omega_0)\big[F(\hat{\theta}_m^1|a) - F(\hat{\theta}_m^1|A)\big] + (\omega_2 - \omega_1)\big[F(\hat{\theta}_m^2|a) - F(\hat{\theta}_m^2|A)\big] \;\approx\; 0.$$

Second, consider males of initial ability b. For $\nu^2 \approx 0$, we have $F(\hat{\theta}_m^1|b) \approx F(\hat{\theta}_m^1|B) \approx 0$, $F(\hat{\theta}_m^2|b) \approx 1$, and $F(\hat{\theta}_m^2|B) \approx 1 - \beta_2/(1-\lambda)$. The expression on the right side of (13) therefore becomes

$$(\omega_1 - \omega_0)\big[F(\hat{\theta}_m^1|b) - F(\hat{\theta}_m^1|B)\big] + (\omega_2 - \omega_1)\big[F(\hat{\theta}_m^2|b) - F(\hat{\theta}_m^2|B)\big] \;\approx\; (\omega_2 - \omega_1)\frac{\beta_2}{1-\lambda}.$$

Third, consider females of initial ability a. For $\nu^2 \approx 0$, we have $F(\hat{\theta}_f^1|a) \approx 1$, $F(\hat{\theta}_f^1|A) \approx \beta_0\lambda - 1$, and $F(\hat{\theta}_f^2|a) \approx F(\hat{\theta}_f^2|A) \approx 1$. The expression on the left side of (14) therefore becomes

$$(\omega_1 - \omega_0)\big[F(\hat{\theta}_f^1|a) - F(\hat{\theta}_f^1|A)\big] + (\omega_2 - \omega_1)\big[F(\hat{\theta}_f^2|a) - F(\hat{\theta}_f^2|A)\big] \;\approx\; (\omega_1 - \omega_0)\Big(2 - \frac{\beta_0}{\lambda}\Big).$$

Finally, consider females of initial ability b. For $\nu^2 \approx 0$, we have $F(\hat{\theta}_f^1|b) \approx F(\hat{\theta}_f^1|B \approx 0$ and $F(\hat{\theta}_f^2|b) \approx F(\hat{\theta}_f^2|B) \approx 1$. The expression on the right side of (14) therefore becomes

$$(\omega_1 - \omega_0)\big[F(\hat{\theta}_f^1|b) - F(\hat{\theta}_f^1|B)\big] + (\omega_2 - \omega_1)\big[F(\hat{\theta}_f^2|b) - F(\hat{\theta}_f^2|B)\big] \;\approx\; 0.$$

If follows that, if $0 < c < \min\{(\omega_1 - \omega_0)(2 - \beta_0/\lambda), (\omega_2 - \omega_1)\beta_2/(1 - \lambda)\}$, all inequalities in (13)–(14) are satisfied for sufficiently small $\nu^2 > 0$. □

References

[1] Aigner, Dennis and Glen Cain. 1977. "Statistical Theories of Discrimination in Labor Markets." *Industrial and Labor Relations Review* **30**, 175–187.

[2] Albrecht, James, Anders Björklundt, and Susan Vroman. 2003. "Is There a Glass Ceiling in Sweden." *Journal of Labor Economics* **21**, 145–177.

[3] Arulampalam, Wiji, Alison Booth, and Mark Bryan. 2007. "Is There a Glass Ceiling over Europe? Exploring the Gender Pay Gap across the Wage Distribution." *Industrial and Labor Relations Review* **60**, 163–186.

[4] Arrow, Kenneth. 1973. "The Theory of Discrimination." In: Orley Ashenfelter and Albert Rees (eds.): *Discrimination in Labor Markets*. Princeton, NJ: Princeton University Press. 3–33.

[5] Becker, Gary. 1957. *The Economics of Discrimination*. Chicago, IL: University of Chicago Press.

[6] Becker, Gary. 1973. "A Theory of Marriage: Part I." *Journal of Political Economy* **81**, 813–846.

[7] Bjerk, David. 2007. "The Differing Nature of Black-White Wage Inequality Across Occupational Sectors." *Journal of Human Resources* **42**, 398–434.

[8] Bjerk, David. 2008. "Glass Ceilings or Sticky Floors? Statistical Discrimination in a Dynamic Model of Promotion and Hiring." *The Economic Journal* **118**, 961–982.

[9] Browne, Kingsley. 1998. "An Evolutionary Account of Women's Workplace Status." *Managerial and Decision Economics* **19**, 427–440.

[10] Browne, Kingsley. 2006. "Evolved Sex Differences and Occupational Segregation." *Journal of Organizational Behavior* **27**, 143–162.

[11] Bureau of Labor Statistics. 2010. *Current Population Survey (Year 2010)* (ftp.bls.gov/pub/special.requests/lf/aa2010/pdf/cpsaat11.pdf).

[12] Chiswick, Barry, Anh Le, and Paul Miller. 2008. "How Immigrants Fare Across the Earnings Distribution in Australia and the United States." *Industrial and Labor Relations Review* **61**, 353–373.

[13] Coate, Stephen and Glenn Loury. 1993. "Will Affirmative-Action Policies Eliminate Negative Stereotypes?" *American Economic Review* **83**, 5, 1220–1240.

[14] de la Rica, Sara, Juan Dolado, and Vanesa Llorens. 2008. "Ceilings or Floors? Gender Wage Gaps by Education in Spain.' *Journal of Population Economics* **21**, 751–776.

[15] del Río, Coral, Carlos Gradín, and Olga Cantó. 2006. "The Measurement of Gender Wage Discrimination: The Distributional Approach Revisited." Society for the Study of Economic Inequality Working Paper 2006-25.

[16] Fang, Hanming and Andrea Moro. 2011. "Theories of Statistical Discrimination and Affirmative Action: A Survey." In Jess Benhabib, Matthew O. Jackson, and Alberto Bisin, editors: *Handbook of Social Economics, Vol. 1A.* The Netherlands: North Holland.

[17] Jenkins, Stephen. 1994. "Earnings Discrimination Measurement: A Distributional Approach." *Journal of Econometrics* **61**, 81–102.

[18] Lundberg, Shelly and Richard Startz. 1983. "Private Discrimination and Social Intervention in Competitive Labor Markets." *American Economic Review* **73**, 340–347.

[19] Milgrom, Paul. 1981. "Good News and Bad News: Representation Theorems and Applications." *Bell Journal of Economics* **12**, 380–391.

[20] Moro, Andrea. 2003. "The Effect of Statistical Discrimination on Black-White Wage Inequality: Estimating a Model with Multiple Equilibria." *International Economic Review* **44**, 467–500.

[21] Phelps, Edmund. 1972. "The Statistical Theory of Racism and Sexism." *American Economic Review* **62**, 659–661.

[22] Rubin, Paul and Chris Paul. 1979. "An Evolutionary Model of Taste for Risk." *Economic Inquiry* **17**, 585–696.

[23] Summers, Lawrence. 2005. Remarks at NBER Conference on Diversifying the Science & Engineering Workforce, January 14, 2005 (www.harvard.edu/president/speeches/summers_2005/nber.php).

[24] USA Today. "Number of female 'Fortune' 500 CEOs at record high." October 26, 2011 (www.usatoday.com/money/companies/management/story/2011-10-26/women-ceos-fortune-500-companies/50933224/1).