Wages and Discrimination with Incomplete Information

Robert E. Marks

Australian Graduate School of Management
University of New South Wales
Sydney, NSW 2052
Australia

bobm@agsm.edu.au

Wages and Discrimination with Incomplete Information

Recent studies of the employer–employee relationship have focussed on the work/shirk decision of the employee, and on incentive-compatible contracts offered by employers, but this paper abstracts from the work/shirk decision to consider non-strategic employees, who nonetheless fall into two categories of productivity: “high” (workers) and “low” (shirkers). The problem facing the wage-taking employer is not to induce working, but to hire employees to maximise profits, when screening of “high” workers is not certain. The paper solves the technical issue of modelling a heterogeneous workforce, and so is able to consider the interaction between existing employees and newly hired labour. This also permits identification of a form of discrimination against a classification of workers believed to be less productive, and discusses how the relative wages of the two classifications may vary as employers’ beliefs of the proportion of “high” in this classification vary.

Keywords: statistical discrimination; heterogeneous labour force; screening; production function.

JEL Classifications: J71, D81, D82

Robert Marks
Australian Graduate School of Management
University of New South Wales
Sydney NSW 2052
Australia

tel: +61-2-9931-9271
fax: +61-2-9313-7279

CONTENTS

LIST OF FIGURES

Wages and Discrimination with Incomplete Information

Robert E. Marks *

1. Introduction.

Does "discrimination" exist in the labour market? Some people have argued that it must, given what they know of employers' behaviour. Others have argued that it does, and still others that it does not, using empirical data as their evidence. It is not the purpose of this paper to argue empirically whether or not discrimination occurs in the labour market, but to further the discussion of the theoretical basis for discriminating behaviour on the part of a profit-maximising employer. The assumption of profit maximising removes the possibility of the employer's discriminating on grounds such as racism, sexism, or mere dislike of people different from him,[1] and means that he will not accept a trade-off between the firm's expected profit and the employment of workers with "adverse" attributes, presumably to maximise his utility. [2] It does not exclude the possibility, as explored by Arrow (1972) and Akerlof (1976), of discrimination which exists partly because of economic incentives.

By "discrimination against" some classification of people we mean a systematic bias against them in the form of lower wages for equal work, even with wage-taking employers. We shall show in Sections 3 and 4 that if there are two types of workers, "highs" and "lows," who are perfect substitutes up to a constant in the production process, and if there are uniform beliefs among risk-neutral employers about the proportions of "highs" in two classifications identified at low cost, then the competitively determined wages of the two classifications will reflect these beliefs, and the workers in the classification believed to be less productive will be paid a wage proportional to this believed shortfall in productivity. So long as the beliefs are true, then, these employees will not on average be discriminated against in the Beckerian sense of being regarded by employers as though their wage were $(1+d)$ times its actual level, where d is Becker's discrimination coefficient. In Section 5, however, we show that the introduction of risk

---

* The author wishes to thank participants at the AGSM's Industrial Organization Workshop, the Fourth Analytical Economists' Workshop, and the 1986 Australasian Meetings of the Econometric Society, at which earlier versions of this paper were presented, in particular Bruce Chapman. The research was partly supported by the Australian Research Council and the Graduate School of Business, Stanford.

Motivation for this paper came from a faculty meeting over twenty years ago in which there was discussion about the desirability of scholarships for women to attend an executive program for senior managers. When the late Malcolm Fisher argued against the proposal, because "there is no evidence in the economics literature of discrimination against women as executives," the late Peter Wilenski (later Australia's Ambassador to the U.N.) responded, "That just reflects on the economics literature." I dedicate this paper to the memories of both men.

1. We shall refer to a male employer throughout, although if she held similar beliefs there would be no reason for a female employer to behave differently.

2. See Becker (1971) and Welch (1967) for models in which discrimination is explained by tastes.

aversion on the part of the employers will result in discrimination, although whether against or towards the less productive classification will depend on the proportions of "highs" and "lows" believed to be in each classification.

The model is a version of "statistical discrimination" (Phelps 1972), where employers possess incomplete information about the productivity of individual workers before hiring them, and base their hiring decisions on easily observed characteristics of the potential employees, such as their membership of a group. See Altonji & Blank (1999) for a survey of economic theories of statistical discrimination. Lundberg & Startz (2004) include statistical discrimination in models of one-sided and two-sided search for trading partners, an elaboration of the labour-market transactions considered here.

The results may also be applied to other economic activities in which there is moral hazard: insurance companies would like to have better information about the riskiness ("quality") of the individual policies which comprise their portfolios; banks when making loan decisions must similarly make judgments about the riskiness of companies as possible defaulters. The results can be generalised to these and other cases of deciding on the hidden quality of individuals on the basis of external, easily observed (or *screened*) characteristics. We are using a non-strategic model of incomplete information, rather than the more popular strategic models which have followed Spence (1973): firms are price-takers in the labour market and workers do not engage in signaling behaviour.

In response to the Darwinian argument that such discrimination is not sustainable, since "under the usual assumption of constant (or increasing) returns to scale, competition would imply the elimination of all but the least discriminatory employers" so that "if there are any non-discriminatory employers, they would drive out the others" (Arrow 1998), we argue that the lower wages of those discriminated against will enable the discriminatory employers to survive: discrimination will lower these wages, and so offset the imperfect information about productivity.[3]

Before demonstrating these results, we discuss in Section 2 how to model a production function incorporating heterogeneous labour, and report some general results about heterogeneous factor inputs, when the production function is non-linear. The technical problem of the marginal product of labour with high- and low-productive labour employed together is solved. Section 3 models the hiring decision of the risk-neutral, profit-maximising firm, and derives a demand curve for labour. Section 4 models the determination of market-clearing wages in both labour markets when the risk-neutral firms have uniform expectations of the productivity of the two classifications. Section 5 models the hiring decisions of risk-averse firms, and identifies a kind of discrimination against the believed less-productive classification. Section 6 relaxes the assumption of uniform expectations. Section 7 places the results in context.

---

3. Farmer & Terrell (1996) proposed a dynamic labour-market discrimination model based on Bayesian updating of beliefs by the employer, who forms initial beliefs of each individual worker's ability conditional on the worker's group, and updates his beliefs each period based on the worker's observed output. If the initial beliefs are incorrect, discrimination may persist despite the learning, they argue.

2. Modelling Heterogeneous Labour.

2.1 Effective labour and the Equal Output property

How should we model heterogeneous labour? For simplicity's sake, let us assume that there are only two types of workers: "low" and "high" productivity. Let the production function of the representative firm be denoted by:

$$Y = G(L, H),$$

where L is the amount of "low" productivity labour and H is the amount of "high" productivity labour, with positive marginal products with respect to the two inputs:

$$G_i(L, H) > 0, \qquad \text{for all } L, H, \ i = L, H.$$

We want to constrain the form of the function $G(.,.)$ to a plausible formulation.

We model heterogeneous labour inputs as

$$Y = F(\, g(\, h\,)\, N\,) \ = \ F(\, H + c\, L\,), \tag{1}$$

where $g(h)$, the *effectiveness factor*, is defined by

$$g(\, h\,) \equiv h + (1 - h\,)\, c,$$

where $h \equiv H/N$ is the proportion of "highs" in the firm's labour force. As discussed in Marks (1993), this formulation exhibits the property of Equal Output,[4] which is based on the assumption that the ratio of the numbers of homogeneous "highs" to "lows" necessary to produce an equal amount of output is constant and less than one, whatever the level of output produced. An example of the *EO* property is the belief (say) that whatever the work to be done, it can be completed by ⅘ the number of redheads as non-redheads, no matter what the size of the task.

Equation (1) models the two types of labour inputs, "high" and "low," as *perfect substitutes up to a constant multiple,* and is a homothetic production function, with parallel isoquants in the $H{-}L$ plane. Marks (1993) proves a theorem which states that only the formulation of equation (1) exhibits perfect substitutability up to a multiple between two types of heterogeneous labour with non-linear production technology.

We can use the effectiveness factor, $g(h)$, to define *effective labour,* $\hat{N}$, as

$$\hat{N} \equiv g(h)\, N.$$

When all workers are "high," $H = N$, $g(h) = g(1) = 1$, and $\hat{N} = N$; and when all workers are "low," $L = N$, $g(h) = g(0) = c$, and $\hat{N} = c\,N$. Equation (1) becomes

$$Y = F(\, g(\, h\,)\, N\,) = F(\, \hat{N}\,) \equiv K(\, h,\, N\,). \tag{2}$$

This formulation parameterises the production function in terms of the measure of effective labour, which is constant along each isoquant. In the $H - L$ plane, these isoquants are straight lines, as implied by homotheticity and the assumption of perfect

4. Marks (1993) also discusses the Equal Numbers property, in which the ratio of maximum levels of output produced by an equal number of "lows" and "highs" alone is constant for any number of workers, and shows that this is identical with the Equal Output property only under constant returns to scale (CRTS).

substitutes up to a multiple.

To determine the marginal product associated with an additional worker, given a specific proportional mix of labour h, we differentiate with respect to number of workers N, to obtain the marginal product of an additional worker,

$$Y_N = \frac{\partial K(h, N)}{\partial N} = g(h)\, F'(N) = [h + (1 - h)c]\, F'(\hat{N}), \tag{3}$$

a linear combination of the two marginal products, $F'(\hat{N})$ and $c\, F'(\hat{N})$. The marginal product of the proportion, h, of "high" workers is given by

$$Y_h = \frac{\partial K(h, N)}{\partial h} = (1 - c)\, N\, F'(\hat{N}). \tag{4}$$

3. The Profit-Maximising Risk-Neutral Employer.

3.1 Imperfect information about workers' productivities

We assume that, as above, the labour force is comprised of two types of worker, "high" and "low" productivity, but that is costly, ex ante, for employers to determine into which of the productivity types a prospective employee falls. Without cost, however, an employer can classify (or screen) a prospective employee into one of two groups, A and B. Moreover, we assume that the employer has the following prior beliefs (or prejudices) about the probability of an employee's belonging to the "high" type: we assume that the employer believes that the probability of an individual from classification A (an "A worker") being a "high" is α, and of being a "low" is $(1 - \alpha)$; we assume that he believes that the probability of an individual from classification B (a "B worker") being a "high" is β, where β is greater than α, and of being a "low" is $(1 - \beta)$, which must be less than $(1 - \alpha)$. That is, we assume that he believes that A workers are less likely to be "high" productivity than are B workers. (This can be thought of as a deterministic belief on the part of the employer about the proportion $h \equiv H/(H + L)$ of "highs" in each of the two classifications, A and B.)

His beliefs are not assumed to be correct. For discrimination to occur it is not necessary that his estimations of the probabilities (or proportions) α and β are correct. We shall discuss below his incentives for revising beliefs and the possible mechanisms for his learning the true probabilities (or proportions).

For the moment, we assume that $\beta = 1$, that is, that the employer believes that all B workers are "highs." From equation (2) the expected output of N individuals who are A workers (with $h = \alpha < 1$) is given by

$$Y(h, N)^A = K(\alpha, N) = F[\hat{N}(\alpha, N)] = F[g(\alpha)\, N],$$

and the expected output of N individuals who are B workers (with $h = \beta = 1$) is given by

$$Y(h, N)^B = K(1, N) = F(N) > Y^A.$$

From equation (3) the marginal product per additional A worker N is given by

$$\frac{\partial Y(h, N)^A}{\partial N} = \frac{\partial K(\alpha, N)}{\partial N} = g(\alpha)\, F'[\hat{N}(\alpha, N)] = [\alpha + (1 - \alpha)c]\, F'[\hat{N}(\alpha, N)],$$

and the marginal product per additional B worker N is given by

$$\frac{\partial Y(h,N)^B}{\partial N} = \frac{\partial K(1,N)}{\partial N} = g(1)\,F'[\hat N(1,N)] = F'[\hat N(1,N)].$$

Hence, the ratio of the two marginal products is given by

$$\frac{Y(\alpha,N)_N^A}{Y(1,N)_N^B} = g(\alpha)\ \frac{F'[\hat N(\alpha,N)]}{F'[\hat N(1,N)]}\,,$$

which in general is not equal to $g(\alpha)$.

Figure 1 depicts the two classifications, A and B, in the $H{-}L$ plane. Since a

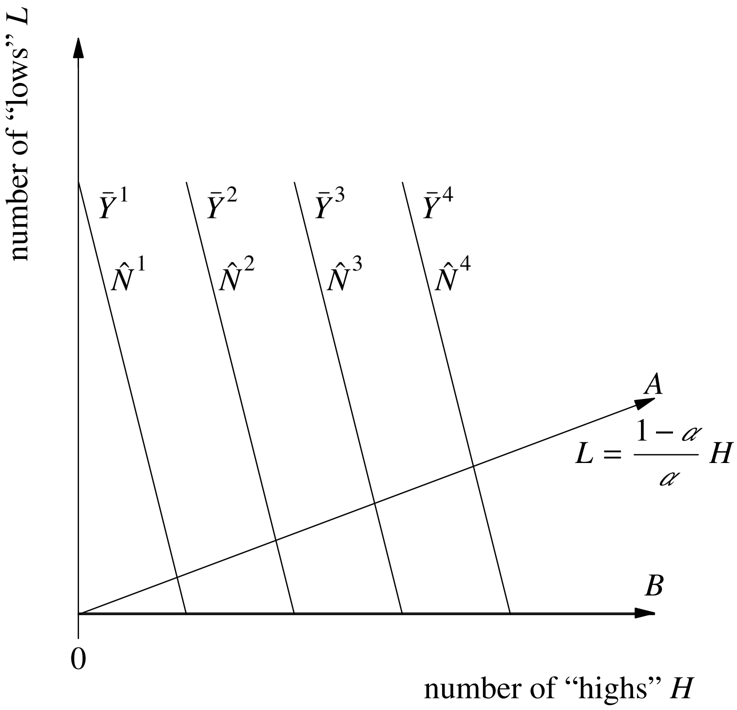

Figure 1. Classifications A and B in the H–L plane.

classification is defined in terms of the mix of "high" and "low" productivity workers, each classification corresponds to a straight line through the origin. Classification B corresponds to the H axis, along which all workers are "highs," and classification A is shown as the ray $O{-}A$, $L = \frac{1-\alpha}{\alpha} H$, with a constant proportion of "highs" and "lows," and with $0 < \alpha < 1$.

3.2 A possible fallacy in hiring

We have assumed that our employer believes that A workers are on average less productive than B workers; we have modelled this belief by assuming that he believes that there is a greater proportion of "low" productivity workers amongst classification A than there is amongst classification B. (In fact, to clarify the algebra, we have assumed that he believes that there are *no* "low" productivity B workers.)

Consider the profit-maximising firm's hiring decision of choosing between two workers: an A worker at wage w^A or a B worker at w^B. One approach is to consider the expected net benefit of one A worker versus the expected net benefit of one B worker. This is incorrect.

To see this, note that the expected net benefit of an i worker is given by

$$\mathrm{E}(NB_i) = g(h^i)\, F' - w^i,$$

where h^i is the proportion of "highs" in classification i. Using the expected net benefit per worker would result in the rule of choosing an A worker if $\mathrm{E}(NB_)$ > $\mathrm{E}(NB_B)$, that is, if

$$g(h^A)\, F' - w^A \ > \ g(h^B)\, F' - w^B,$$

or, rearranging, if

$$\frac{F' - w^A / g(h^A)}{F' - w^B / g(h^B)} \ > \ \frac{g(h^B)}{g(h^A)}\,.$$

Let us define v^i, the *effective wage,* as the ratio of the wage to the effectiveness factor for classification i:

$$v^i \equiv w^i / g(h^i).$$

This results in the same wage bill in actual or effective labour, since $w\, N = v\, \hat{N}$.

If there are equal effective wages for the two classifications, $v^A = v^B$, then choose from A if $g(h^A) > g(h^B)$. But by definition this cannot be, so even with equal effective wages the choice will be to choose from B. Indeed, even if

$$w^A \ < \ \frac{g(h^A)}{g(h^B)}\, w^B \ < \ w^B,$$

then choose from B.

This is incorrect: the fallacy is that by looking at the expected net benefit *per worker* we have biased our decision rule in favour of the more productive: even if the ratio of the lower classification's wage to the upper classification's wage is proportional to the ratio of their respective (believed) productivities ($w^A / w^B = g(h^A) / g(h^B)$), using this choice rule the employer would choose the more productive. This is because the expected net benefit *per worker* will be greater for the more productive.

This error is similar to the mistake in cost-benefit analysis of ranking alternative projects by their net present values and then choosing projects with successively smaller net present values, until the capital budget constraint becomes binding on the decision-maker. The correct approach is to rank by the ratios of the gross present values of benefits to costs. In the hiring decision, this would be equivalent to comparing not expected net benefit per worker but expected net benefit per wage dollar spent.

Although a moment's thought shows the fallacy of comparing classifications of workers by expected net benefit per worker, we note that making this comparison would result in discrimination against workers in the classification believed to be of lower productivity *even with risk neutrality:* if an employer believed that A workers were half as productive as B workers but commanded a wage of only half, then—so long as there was CRTS[5] (so that two A workers produced as much as one B worker)—the employer should be indifferent between hiring two A workers and one B worker. If, however, he only considered the expected net benefit per worker rather than the per wage dollar spent, he would prefer a B worker to an A worker. Has this error in thought been responsible for much actual discrimination?

3.3 The hiring decision

Consider the decision of a risk-neutral employer to hire an additional worker. If the additional worker is an A, the employer will expect less additional output than if the additional worker is a B. If the wages of the two classifications were equal, we should expect the employer to prefer to hire from classification B, even though the employer believed that some proportion α of A workers were "highs." But what if the wage for an A worker were lower than the wage for a B worker? We might expect a profit-maximising employer to prefer hiring from classification A if the ratio of marginal output to wage were greater for A than for B. What is the ratio of marginal products?

Theorem 1: *At any combination of workers from classifications A and B the expected marginal product of an additional A worker is a fraction of the marginal product of an additional B worker, where the fraction is equal to $g(\alpha)$, the effectiveness factor of classification A, independent of the returns to scale.*[6]

This result is independent of the returns to scale of the homothetic production function, and independent of past hirings: the expected marginal product of an additional A worker is *always* $g(\alpha)$ times the expected marginal product of an additional B worker. Moreover, if the employer's belief of the proportion of "highs" in classification A changed, so would his expectation of the marginal product of an A worker, as given by equation (A1). With static, homogeneous expectations of the effectiveness factors of classifications, and with risk-neutral employers, we would lose nothing in generality by considering a classification A of all "lows" as well as our classification B of all "highs," but we persist with our "$\alpha > 0$" classification A, at some notational cost, in order to model risk-averse employers, heterogeneous beliefs across employers, and changes in belief of the value of classification A's proportion of "highs," α.

From Theorem 1, so long as the ratio of classification A wages to classification B wages equals the ratio of effectiveness factors, the risk-neutral employer will be

---

5. With diminishing returns to scale (DRTS), we run into problems using marginal analysis. At first blush we say that the second A worker will not produce as much as the first (which must follow if there is DRTS) in which case two A workers will not equal one B worker. To model this we must use finite increments, not the infinitesimal increments of marginal analysis, so that the incremental product is the slope of the chord rather than the tangent of the total product curve plotted against labour input. / /

6. See the Appendix for all proofs.

indifferent between classifications A and B as a source for additional workers. Before it is possible to examine how the wage rates for the two classifications might be competitively determined, and with what outcomes, given heterogeneous expectations of the effectiveness factors, we must first consider the optimum, profit-maximising levels of employment by the firm, and its maximum profit, assuming DRTS.

The risk-neutral, profit-maximising firm faces two decisions: which classification to hire from, and how much labour to employ. The second problem can be written as

$$\max_{N^i} \ \pi(N^i) = F[\, g(h^i)\, N^i] - w^i N^i, \quad F' > 0,\ F'' < 0, \tag{5}$$

where π is the expected profit and w^i is the wage rate paid to classification i. With DRTS, necessary and sufficient conditions for maximising expected profit are that the marginal product equals the wage:

$$g(h^i)\, F'[\, g(h^i)\, N^{i}*] = w^i. \tag{6}$$

From the definition of effective wage, equation (6) states that the risk-neutral, profit-maximising employer will choose that level of employment of classification i workers that equates the marginal product of effective labour, $g(h^i)\, N^i$, with the effective wage, v^i. From equation (6) the optimum derived demand for classification i labour, $N^i\,*$, can be determined as a function of the effective wage rate, v^i:

$$N[w^i, g(h^i)]\,* = \tfrac{1}{g(h^i)}\ F'^{-1}(v^i). \tag{7}$$

Substituting for $N^i\,*$ into equation (5), we obtain the maximum expected profit, $\pi^i\,*$, from employing classification i workers:

$$\pi^i\,* = \pi^*\,(v^i) = F\Big[F'^{-1}(v^i)\Big] - v^i F'^{-1}(v^i). \tag{8}$$

Differentiating this maximum expected profit with respect to the effective wage v^i, we obtain

$$\begin{aligned} \frac{d\pi^i\,*}{dv^i} &= \left(F'\Big[g(h^i)\, N^i\,*\Big] - v^i\right)\frac{1}{F''} - F'^{-1}(v^i) \\ &= -\, g(h^i)\, N^i\,* \\ &= -\, \hat{N}(h^i, N^i*) < 0. \end{aligned} \tag{9}$$

Theorem 2: *The risk-neutral, profit-maximising firm can separate its hiring decisions into two: first, hiring from the classification with the lowest effective wage, and, second, hiring until the marginal product of effective labour equals the effective wage.*

Corollary: *When the effective wage is equal across classifications, the total wage bill associated with risk-neutral, profit-maximising employment is also equal across classifications.*

3.4 A specific production function

A specific production function which exhibits decreasing returns to scale is

$$Y(h, N) = F[g(h)\,N] = f\,[g(h)\,N]^{\lambda}, \qquad 0 < \lambda < 1, \tag{10}$$

where f is a constant. The parameter λ is equal to the share of labour in total output if the marginal product of labour equals the wage. Substituting equation (10) into equation (6), we get

$$\pi^i * = \frac{1-\lambda}{\lambda} (\lambda f)^{\frac{1}{1-\lambda}} (v^i)^{-\frac{\lambda}{1-\lambda}} \tag{11}$$

and

$$N^i * = \left(\lambda f \, \frac{g(h^i)^{\lambda}}{w^i}\right)^{\frac{1}{1-\lambda}} = \frac{1}{g(h^i)} \left[\lambda f v^i\right]^{\frac{1}{1-\lambda}}. \tag{12}$$

Hence, we can conclude that, for the DRTS production function of equation (10), so long as the effective wage paid to classification i, v^i, is equal across classifications, the risk-neutral employer will be indifferent as to hiring from any of the classifications i, since his expected profit will be equal, from equation (11).

From equation (12) the optimum workforce N * is an increasing function of the effectiveness factor $g(h)$, but a decreasing function of the wage rate, ceteris paribus. The relationship between the optimum workforce and the wage rate is plotted in Figure 2, a hyperbola.

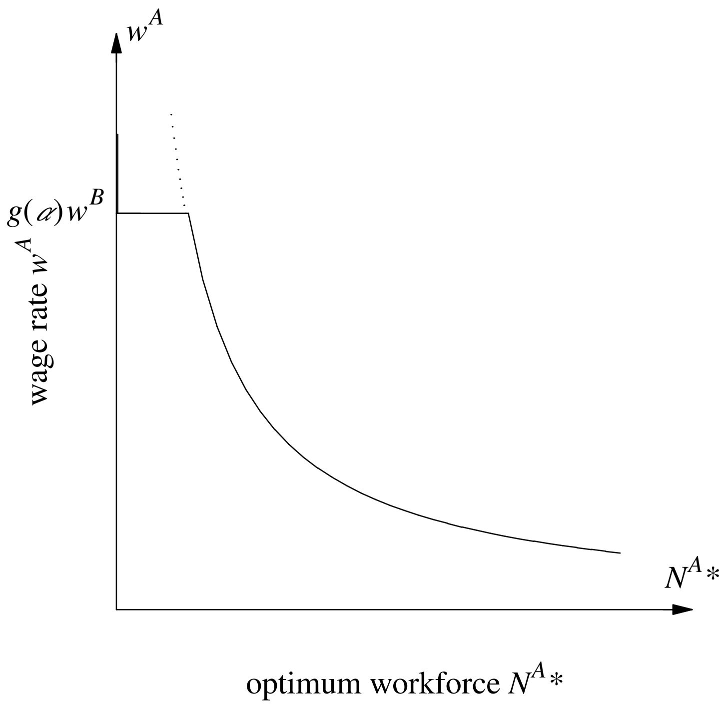

Figure 2. The derived demand for labour.

Equation (12) does not, however, admit of the possibility of substitution in hiring between one classification and another as the relative wages change. From Theorem 2, so long as the effective wage of classification A, v^A, is less than the effective wage of classification B, v^B, then the risk-neutral employer will hire from A, but when v^A is greater than w^B, from B. From the definition of the effective wage and the effectiveness factor, the effective wage is equal to the actual wage for classification B, with $h = \beta = 1$:

$$v^B = \frac{w^B}{g(1)} = w^B.$$

The drop in demand for A workers above a threshold wage is shown in Figure 2 as a truncation to zero of the level of N^A for $w^A > g(\alpha)\, w^B$, as substitution occurs to B workers, for constant $g(\alpha)$ and w^B.

We have modelled a risk-neutral, profit-maximising firm facing two classifications of workers, B of all "highs," and A with a proportion $\alpha < 1$ of "highs." The production technology was discussed in Section 2 above. The firm will chose to hire from the classification with the lower effective wage, and will hire up to the point at which the marginal product of effective labour equals the effective wage. In the sections below we consider the competitive determination of the relative wages of the two classifications with uniform expectations of the proportions of "highs."

4. Uniform Expectations.

Aggregating across risk-neutral, profit-maximising firms with uniform expectations of the proportion of "highs" in each classification and with uniform production technology, we can model the market demand for A workers, $D(w^A)$, as shown in Figure 3. When the wage for A workers, w^A, is less than $g(\alpha)\, w^B$, then the demand for these workers is a decreasing function of the wage; when the wage is greater, then the demand is zero.

Because of the character of the decision, there is no shift in the demand curve for A workers as the wage for B workers falls: the line of truncation, LM, moves down instead. Figure 3 depicts two possible supply schedules for A workers: the market-clearing wage corresponding to supply ${}^1S^A$ is ${}^1w^A$ at a quantity of ${}^1N^A$ workers; the market-clearing wage corresponding to the tighter supply of ${}^2S^A$ is given by ${}^2w^A = g(\alpha)\, w^B$ and the quantity is ${}^2N^A$.

Will arbitrage in the two markets—those for A workers and for B workers—always result in equal effective wages for the two classifications; that is, will $v^A = v^B$ in equilibrium ?

Theorem 3: *If the belief that a classification i worker produces the fraction $g(h^i)$ of the output of a "high" worker is uniformly held by all employers, then the competitive determination of the wages of all classifications with voluntary choice on the part of risk-neutral employers will result in equal effective wages in all classifications.*

We need to establish that the equilibrium is stable, and to explore the quantities of labour hired. Theorems 2 and 3 and the linearity of the model allow us to plot the supply and demand schedules on one graph, Figure 4, in the effective labour–effective wage plane. As we have discussed in Section 3, the effectiveness factor of classification B, $g(h^B)$, is equal to one, since all B workers are thought to be "highs," and the effectiveness

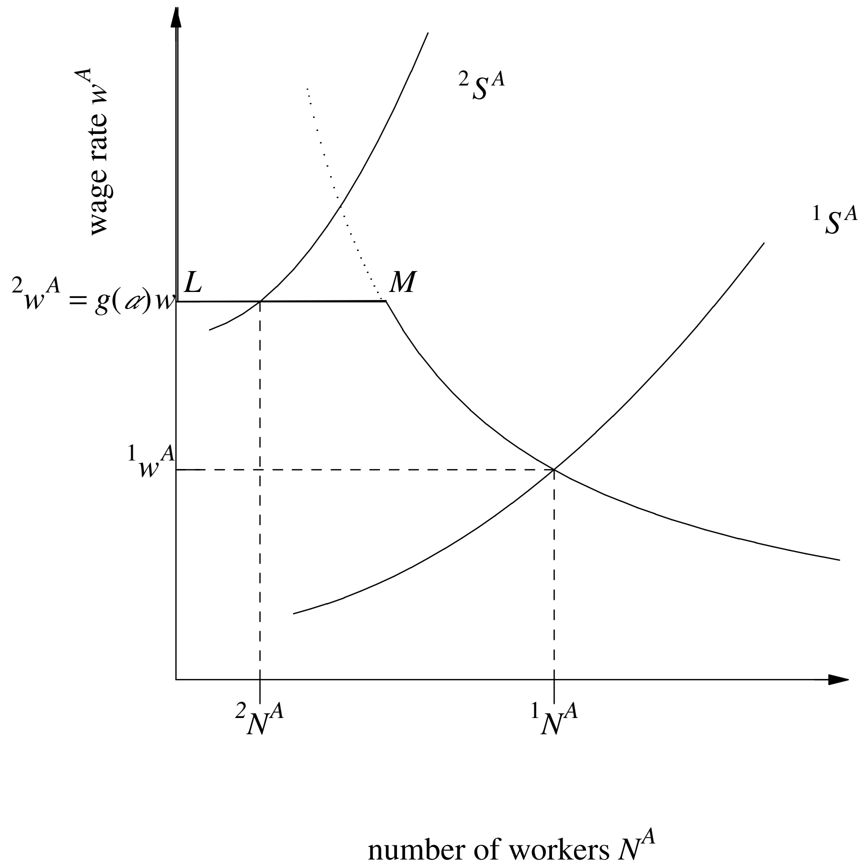

Figure 3. Two possible supply schedules for A workers.

factor of classification A is given by $g(\alpha)$, where there are believed to be a proportion α "highs" in classification A. Hence, from equation (A1), effective labour in classification B is given by numbers of B workers, N^B, and in classification A by $g(\alpha)\, N^A$; and the effective wage in classification B by the wage, w^B, and in classification A by $\frac{1}{g(\alpha)}\ w^A$.

In Figure 4 we have plotted the derived demand for effective labour as a function of effective wage, which is equal to the derived demand for B labour as a function of the B wage, $D^B(w^B)$. We have horizontally summed the supply of effective B labour, $S^B(w^B)$ and the supply of effective A labour, $g(\alpha)\, S^A[\frac{w^A}{g(\alpha)}]$.

Stable, market-clearing equilibrium occurs at wage $\bar{w}^B$ and quantity of effective labour $\overline{\hat{N}}$, given by

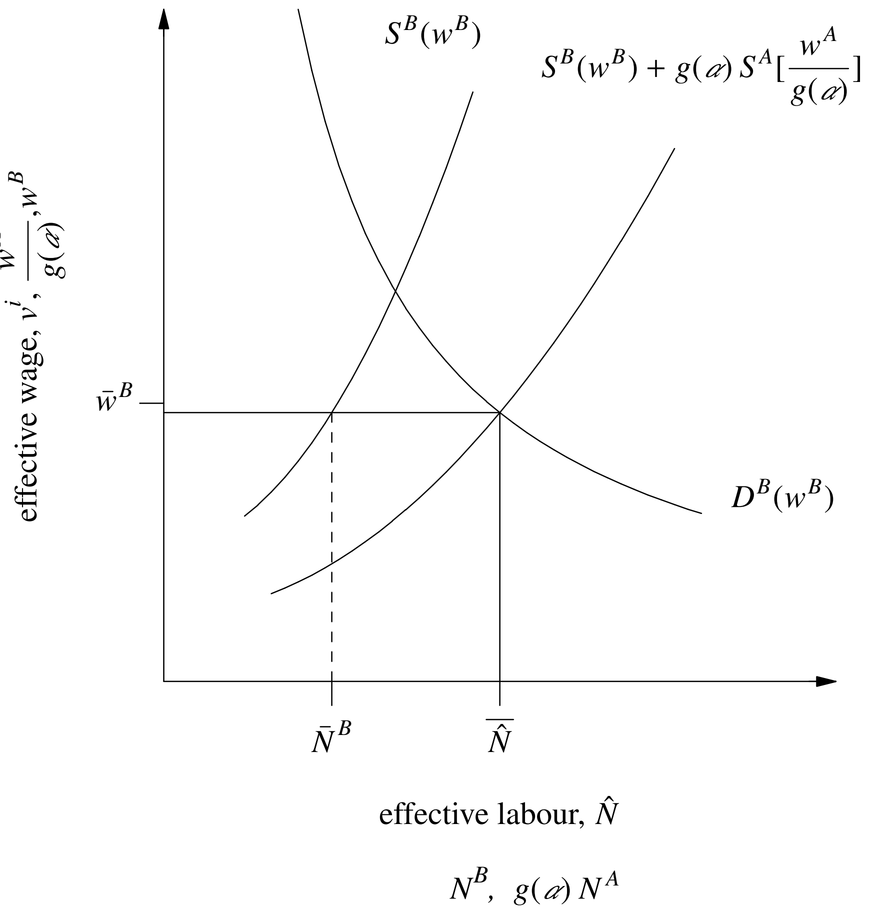

Figure 4. The market for effective labour.

$$\overline{\hat{N}} = \bar{N}^B + g(\alpha)\, \bar{N}^B = D^B(\bar{w}^B), \tag{13}$$

where, as shown in Figure 4,

$$\bar{N}^B = S^B(\bar{w}^B),$$
$$\bar{N}^A = S^A[\tfrac{\bar{w}^A}{g(\alpha)}] = S^A(\bar{w}^B).$$

How would a change in employers' belief in α, the proportion of "highs" believed to occur in classification A, affect the competitively determined effective wage, $\bar{w}^B$? From its definition,

$$\frac{\partial g}{\partial \alpha} = 1 - c > 0,$$

and so as α increases, the effectiveness factor also increases. Differentiating equation (13) with respect to α, and noting that, from Theorem 3, $\bar{w}^B = \frac{\bar{w}^A}{g(\alpha)}$, we obtain

$$\frac{d\bar{w}^B}{d\alpha} = \frac{(1-c)\,S^A(\bar{w}^B)}{D^{B\prime} - S^{B\prime} - g(\alpha)\,S^{A\prime}} < 0.$$

Intuitively, this makes sense: as the proportion of "highs" believed to occur in classification A increases, it is as though the supply of labour increases, thus driving down the wage.

5. Risk-Averse Employers.

Theorem 4: *Even if the effective wages of the two classifications' workers are equal, the risk-averse, utility-maximising firm will still prefer to hire from a classification believed to comprise wholly high-productivity workers and not from a classification believed to comprise a mix of high- and low-productivity workers.*

Corollary: *The wage of the classification of heterogeneous workers must be "disproportionately" lower than their believed expected productivity before risk-averse employers will hire from this classification rather then the wholly high classification.*

Another way of putting this is to say that the employer acts as if the effective A wage is $v^A(1+d)$, when in fact the effective wage is only v^A. That is, risk aversion introduces discrimination which can be described in terms similar to Becker's (1971) discrimination coefficient, discussed below.

Thus, in this simple model, risk aversion provides a possible explanation of discrimination against workers from classification A: even if the relative wages commanded by the two classifications accurately reflect the proportion of low-productivity workers in the mixed classification, risk-averse employers will still prefer to hire from the wholly high-productivity classification. Of course, an under-estimate of the proportion of high-productivity workers in the cheaper classification will lead to further discrimination.

The proof of Theorem 4 relied on the homogeneity of classification B: $\beta = 1$ and $g(h^B) = 1$. Can we derive a more general result if there is (believed) heterogeneity of workers in both classifications? Unfortunately, using the common Rothschild & Stiglitz (1970) comparison, we cannot order the two classifications.

Following Rothschild & Stiglitz, we define the two cumulative distribution functions, $A(x)$ and $B(x)$, where

$$A(x) \equiv \text{Prob}\ \{\ X \le x\},$$

which results in

$$A(x) = \begin{cases} 0 & x < cF'/w^A \\ 1-\alpha & cF'/w^A \le x < F'/w^A \\ 1 & F'/w^A \le x \end{cases}$$

and a corresponding function for $B(x)$. According to Rothschild & Stiglitz, two conditions are necessary and sufficient that distribution B is "less risky" than distribution A: first, that the expected values of the two distributions are equal, and second that

$$\int_0^y [A(x) - B(x)]\, dx \geq 0, \qquad \text{for all } y \in [0,\, F'/w^A]. \tag{14}$$

The first condition can be written as

$$\alpha \frac{F'}{w^A} + (1-\alpha)\frac{cF'}{w^A} = \beta \frac{F'}{w^B} + (1-\beta)\frac{cF'}{w^B}\,,$$

where by definition $\alpha < \beta$ and $w^A < w^B$. This is equivalent to equal effective wages across the two classifications:

$$\frac{w^A}{g(h^A)} = \frac{w^B}{g(h^B)}\,.$$

To understand the second, consider Figure 5, which plots the two distributions $A(x)$ and $B(x)$ against marginal product per wage dollar spent, where the believed proportion of "highs" in classification A is α, which is less than β for classification B. The condition requires that up to any point z the area under the curve $A(z) - B(z)$ always be non-negative. From the figure it is readily seen that for *any* uncertainty associated with classification B (that is, for $\beta < 1$) the condition of equation (14) will not be satisfied for the range $cF'/w^B < z < cF'/w^A$, in which $A(z) = 0 < B(z)$.

That the Rothschild-Stiglitz conditions do not hold means that we cannot say which of the distributions A and B is the riskier, that is, we cannot say which classification would be preferred by *all* risk-averse employers, so long as both classifications are believed to be heterogeneous. (For von Neuman-Morgenstern utility maximisers, higher variance is not sufficient for lower utility.) It is possible that the Diamond & Stiglitz (1974) partial ordering would shed more light on a comparison of distributions A and B relative to a specific utility function, but we leave such analysis to a later paper.

6. Non-Uniform Expectations.

Relaxation of the assumption of uniform expectations will result in distributions of the proportion of "highs" in each classification across employers. We can define the effective wage to employer j for classification i, v^i_j, as

$$v^i_j \equiv \frac{w^i}{h^i_j + (1-h^i_j)c}\,,$$
$$= w^i / g(h^i_j),$$

where h^i_j is the proportion of "highs" believed in classification i by employer j. Then, from Theorem 2, a risk-neutral employer j will hire from classification A if the effective A wage is less than the effective B wage, $v^A_j < v^B_j$. Aggregating across employers will result in market demands for A and B workers, and market clearing in the two markets will result in the actual wages w^A and w^B, which will be functions of the distributions of beliefs, h^A_j and h^B_j across employers. We leave for a later paper a thorough investigation of the equilibrium wages.

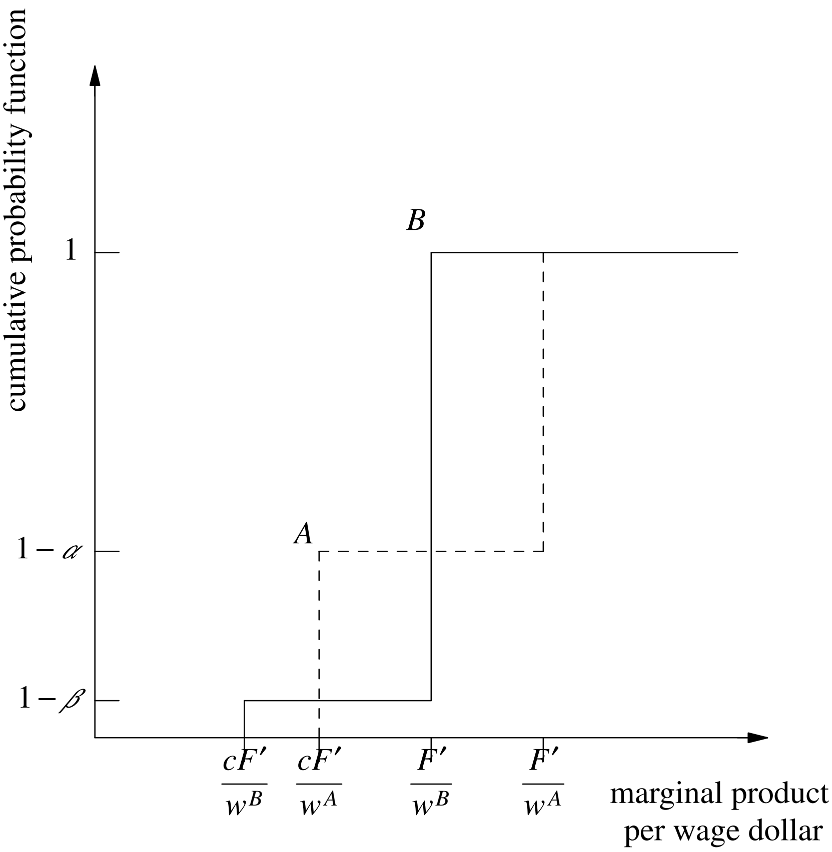

Figure 5. The two distributions, $A(x)$ and $B(x)$.

7. Discussion and Conclusions.

Becker (1971), in his pioneering study, defines his *discrimination coefficient* as follows: if the wage of the discriminated-against classification = w, but the employer acts as if the wage is $w(1 + d)$, then d is the discrimination coefficient. Further, he defines his *market discrimination coefficient* as

$$\frac{w^{disc} - w^{non-disc}}{w^{disc},}$$

where w^{disc} is the market wage of the discriminated, and $w^{non-disc}$ is the market wage of the non-discriminated. He shows that perfect substitutability results in the individual discrimination coefficient equalling the market discrimination coefficient.

What we have shown is that in a model with competitively determined wages and uniformly held (if not necessarily correct) beliefs about the productivity of a classification of workers vis-à-vis other classifications of workers, then with risk-neutral employers there will be no discrimination in Becker's sense: $d = 0$, and the wage differentials between classifications will reflect the universal belief in the relative differences in productivity between the specific classification and other classifications. Becker's market discrimination coefficient will equal $1 - g(\alpha)$ in our terms, where $g(h)$ is the effectiveness factor, as defined in Section 3. Risk-averse employers, however, will behave as though Becker's $d \neq 0$ for effective wages, and will hire from the less risky classification, ceteris paribus.

Becker is concerned with exploring what might be called non-profit-maximising discrimination, in which employers are willing to incur the cost of lower profits presumably to maximise their utilities, whereas we have modelled employers who are maximising their profits, subject to their (perhaps incorrect) information about the proportions of "high" and "low" productivity workers in easily identified classifications. Other studies, Arrow (1972, 1973), have abstracted from differences in productivity between classifications of workers, which is equivalent to *correctly* setting $h^i = h^j$ for all i, j in our model, where h^i is the proportion of "highs" in classification i, as defined in Section 3.

Arrow (1973) discusses a model of imperfect information on the part of employers similar to ours. He concludes that discrimination can result not from employers' tastes but from their perceptions of reality if three conditions hold: first, the employer must be able to distinguish A workers from B workers (in our terms) at a reasonably low cost; second, the employer must incur some cost before he can determine the employee's true productivity (whether he is a "high" or a "low" in our model); and, third, that the employer has some preconception of the distribution of productivity across each of the two classifications of workers (our $h^A = \alpha$, and $h^B = 1$). We have assumed the first and discussed the third at length. The second condition is to provide a reason for the employer to use the "surrogate" information on productivity provided by the classification membership: otherwise the employer could simply fire the "lows," or pay them according to their productivity.

If an employer never hired A workers, because, say, he had to pay them the full wage but anticipated getting less than full productivity from them on average, then he would never have the opportunity to learn their true productivity. (Employment of a worker might not, of course, be sufficient to convince the employer that he was wrong in his belief.) In the face of a wide-spread belief that A workers were less productive, however, an "affirmative action" programme (Coate & Lowry 1983), apart from any equity justifications, might be advocated on the grounds that it encouraged employers to face their beliefs, and perhaps to learn whether they were prejudiced. A competitive environment for the employers' products might also result in the realisation that competitors were gaining an edge by hiring A workers, who were not in practice as unproductive as previously believed.

A paper by Lewis (1979), however, presents empirical evidence that attitudes will not be easily altered by evidence of error. Lewis argues that employers misperceive the relative quit rates of male and female workers; they are aware that, in general, women have higher quit rates but they exaggerate the difference, despite the experience of their own companies, by under-estimating the quit rates of men and over-estimating the quit rates of women, on average and to a significant degree. To the extent that employers act on their inaccurate assessments, then they will discriminate against women when making hiring, training, and promoting decisions.

Phelps (1972) introduced the notion of "statistical discrimination" with a model in which, sampling from a population of job applicants, each employer is able to measure each applicant's productive potential with a test score which is a statistical measure of the applicant's potential productivity. Our model includes greater variability in productive potential for *A* applicants than for *B* applicants but no test scores. Consequently, it does not include any of Phelps' cases. In examining the choice of a risk-averse employer as the variability of potential changes, however, our model focuses on statistical discrimination in a different way from Phelps, by including the wage rates to be paid to workers from the two classifications.

Arrow (1973) examines the possible basis for employers' believing that $h^A \neq h^B$ even when there is no intrinsic difference in abilities between *A* and *B* workers, and concludes that further work is needed on how individual employers acquire knowledge which will modify their initial estimates of distributions of "highs" and "lows" across classifications, on how these perceptions affect the market for labour, and hence on any incentives to hire workers, and even for workers to improve their individual ability and performance.

The next step in the research programme will be to parameterise the heterogeneous pattern of belief of the effectiveness factor across employers. Preliminary work suggests, as one would expect, that the results change considerably, as employers are faced with market relativities in the wages that do not accord with their individual beliefs in the relative productivity of workers from the classifications.

Appendix

Theorem 1: At any combination of workers from classifications *A* and *B* the expected marginal product of an additional *A* worker is a fraction of the marginal product of an additional *B* worker, where the fraction is equal to $g(\alpha)$, the effectiveness factor of classification *A*, independent of the returns to scale.

Proof: Consider the specific combination of $\bar{H}$ "high" productivity and $\bar{L}$ "low" productivity workers; this can also be described as $(\bar{h}, \bar{N})$. From equations (3) and (4) the expected marginal product of an additional worker from classification *i*, where $i = A$ or B, can be calculated as

$$\begin{aligned} \frac{dY(\bar{h},\bar{N})}{dN^i} &= F'[\hat{N}(\bar{h},\bar{N})]\left(\frac{\partial \hat{N}}{\partial h}\frac{dh}{dN^i} + \frac{\partial \hat{N}}{\partial N}\frac{dN}{dN^i}\right) \\ &= F'[\hat{N}(\bar{h},\bar{N})]\left((1-c)\bar{N}\,\frac{h^i-\bar{h}}{\bar{N}} + g(\bar{h})\right) \\ &= F'[\hat{N}(\bar{h},\bar{N})]\,[h^i + (1-h^i)c] \\ &= F'[\hat{N}(\bar{h},\bar{N})]\,g(h^i) \\ &= h^i F' + (1-h^i)\,cF'. \end{aligned} \tag{A1}$$

where h^i is the proportion of "highs" associated with classification *i*: h^A is α and h^B is one. Thus the ratio of the marginal product of an additional *A* worker to the marginal product of an additional *B* worker is given by

$$\frac{dY(\bar{h},\bar{N})/dN^A}{dY(\bar{h},\bar{N})/dN^B} = \frac{g(\alpha)}{g(1)}\,. \tag{A2}$$

But, from the definition of the effectiveness factor, $g(1) = 1$, so the ratio of equation (A2) is equal to $g(\mathcal{A})$, the effectiveness factor of classification A. □

Theorem 2: The risk-neutral, profit-maximising firm can separate its hiring decisions into two: first, hiring from the classification with the lowest effective wage, and, second, hiring until the marginal product of effective labour equals the effective wage.

Proof: From equation (6), with workers from any classification the expected profit is maximised when the marginal product of effective labour equals the effective wage for that classification. Equation (8) demonstrates that the maximum expected profit is a function of the effective wage, and equation (9) demonstrates that this is an inverse relationship: the lower the effective wage, the higher the maximum expected profit. Theorem 1 allows us to generalise from two to many classifications. Thus the risk-neutral firm will choose to hire from the classification with the lowest effective wage, and to hire up to the point where the marginal product of effective labour equals this wage. □

Corollary: When the effective wage is equal across classifications, the total wage bill associated with risk-neutral, profit-maximising employment is also equal across classifications.

Proof: From equation (7), at risk-neutral, profit-maximising levels of employment the effective labour levels, $g(h^i)\, N^{i\,*}$, are equal across classifications if the effective wage, v^i, is equal across classifications. But in this case, from the definition of effective wage, the effectiveness factor for each classification, $g(h^i)$, is proportional to the wage, w^i. Hence, the total wage bill for each classification, $w^i N^{i\,*}$, is equal across classifications. (This also follows from the linearity of the problem.) □

Theorem 3: If the belief that a classification i worker produces the fraction $g(h^i)$ of the output of a "high" worker is uniformly held by all employers, then the competitive determination of the wages of all classifications with voluntary choice on the part of risk-neutral employers will result in equal effective wages in all classifications.

Proof: From Theorem 2 the risk-neutral, profit-maximising firm will hire from the classification with the lowest effective wage. With competitively determined wages the wage for classification i, w^i, cannot be greater than the equivalent wage of any other classification j

$$w^i \le \frac{g(h^i)}{g(h^j)}\, w^j, \qquad \text{for all } i, j.$$

But if

$$w^i < \frac{g(h^i)}{g(h^j)}\, w^j, \qquad \text{for any } j,$$

it follows that

$$w^j > \frac{g(h^j)}{g(h^i)}\, w^i,$$

which cannot hold, by the assumption of competitively determined prices. Thus it must follow that

$$w^i = \frac{g(h^i)}{g(h^j)}\, w^j, \qquad \text{for all } i, j,$$

or

$$v^i = v^j, \qquad \text{for all } i, j.$$

From Theorem 2, the associated quantity will be N^i if exclusive hiring from classification i, where

$$N^i = \frac{g(h^j)}{g(h^i)}\, N^j, \quad \text{for all} \;\; i, j. \;\square$$

Theorem 4: Even if the effective wages of the two classifications' workers are equal, the risk-averse, utility-maximising firm will still prefer to hire from a classification believed to comprise wholly high-productivity workers and not from a classification believed to comprise a mix of high- and low-productivity workers,

Proof: We assume that the risk-averse firm is an expected-utility maximizer, where the utility function $U(.)$ is positive and increasing in profit, but strictly concave, so that

$$\alpha U(X) + (1-\alpha)U(Y) \;<\; U[\alpha X + (1-\alpha)Y], \qquad 0 \le \alpha \le 1. \tag{A3}$$

Let the wholly high-productivity-worker classification be B, and the mixed classification be A. The marginal product of high-productivity labour is F', and of low-productivity labour cF', where $0 < c < 1$.

The firm will choose from classification B if the expected utility per wage dollar spent on hiring B labour is greater than the expected utility per wage dollar spent on hiring A labour. That is, the firm will choose from B if

$$\alpha U\left(\frac{F'}{w^A}\right) + (1-\alpha)U\left(\frac{cF'}{w^A}\right) \;<\; U\left(\frac{F'}{w^B}\right), \tag{A4}$$

where α is the believed proportion of high-productivity workers in A, and $(1-\alpha)$ the believed proportion of low. All B workers are believed to be high-productivity.

Assume that there are equal effective wages for each classification ($v^A = v^B$), or that

$$\frac{g(h^A)}{w^A} = \frac{\alpha + c(1-\alpha)}{w^A} = \frac{1}{w^B} = \frac{g(h^B)}{w^B},$$

so that the effective wage is the same for workers from both classifications. Then, from (A4), the firm will choose from B if

$$\alpha U\left(\frac{F'}{w^A}\right) + (1-\alpha)U\left(\frac{cF'}{w^A}\right) \;<\; U\left\{[\alpha + (1-\alpha)c]\,\frac{F'}{w^A}\right\},$$

which follows from the concave, risk-averse utility function of (A3). Hence, even if the effective wages of A and B workers are equal, the risk-averse firm will hire from B. $\square$

Corollary: The wage of the classification of heterogeneous workers must be "disproportionately" lower than their believed expected productivity before risk-averse employers will hire from this classification rather then the wholly high classification.

Proof: Since firms are risk averse, they will forgo the higher marginal product per wage dollar spent from hiring A workers when the effective wages are equal, in order to reduce their risk about the type of workers they are hiring. $\square$

References:

Akerlof, George "The economics of caste and the rate race and other woeful tales." *Quarterly Journal of Economics,* **90** (1976): 599–61.

Altonji, Joseph G., and Blank, Rebecca M., "Race and Gender in the Labor Market," in Orley Ashenfelter and David Card, eds. *Handbook of Labor Economics, Volume 3C,* Amsterdam: Elsevier Science, 1999, 3143–3260.

Arrow, Kenneth J. "Some mathematical models of race discrimination in the labour market." In *Racial Discrimination in Economic Life,* edited by A.H. Pascal. Lexington: Lexington Books, 1972.

_________. "The theory of discrimination." In *Discrimination in Labor Markets,* edited by O. Ashenfelter and A. Rees. Princeton: Princeton University Press, 1973.

_________. "What Has Economics to Say about Racial Discrimination?" *The Journal of Economic Perspectives*, **12**(2): 91–100, 1998.

Becker, Gary. *The Economics of Wage Discrimination.* Chicago: University of Chicago Press, second edition, 1971.

Coate, Stephen, and Glenn Loury, "Will Affirmative Action Policies Eliminate Negative Stereotypes?" *American Economic Review*, **83**, 1220–42, 1983.

Diamond, Peter A., and Stiglitz, Joseph E. "Increases in risk and in risk aversion." *Journal of Economic Theory* **8** (1974): 337–60.

Farmer, Amy and Terrell, Dek, Discrimination, Bayesian updating of employer beliefs, and human capital accumulation, *Economic Inquiry,* **34**(2): 204–219, 1996.

Lewis, D. E. "Comparative quit rates of men and women." *Journal of Industrial Relations* **21** (1979): 331–50.

Lundahl M., and Wadensjö E. *Unequal Treatment: A Study in the Neo-Classical Theory of Discrimination.* New York: New York University Press, 1984.

Lundberg, Shelly, and Startz, Richard, Information and Racial Exclusion, November 2004. http://www.econ.washington.edu/user/startz/Working_Papers/Race.pdf

Marks, Robert E. "Modelling heterogeneous inputs," *Economics Letters,* **42**(2–3): 167–171, 1993.

Phelps, Edmund S. "The statistical theory of racism and sexism." *American Economic Review* **62** (1972): 659–62.

Rothschild, Michael, and Stiglitz, Joseph E. "Increasing risk: 1. A definition." *Journal of Economic Theory* **2** (1970): 225–43.

Spence, Michael "Job market signaling." *Quarterly Journal of Economics* **87** (1973): 355–379.

Welch, F. "Labor-market discrimination: extrapolation of income differences in the rural South." *Journal of Political Economy,* **75** (1967): 584–596.