Promotion-Driven Inequality*

Daniel Ferreira Radoslawa Nikolowa Elena S. Pikulina

February 14, 2025

Abstract

We present a model of promotions where vacancies, wages, and internal career paths are endogenously determined in a labor market equilibrium. Workers differ only by a non-productive label: "blue" or "red." Firms are biased towards blue workers in promotions. In equilibrium, two types of firms emerge: (i) "Mixed firms" hire both types of workers, offer risky career paths (where promotion is not guaranteed), are large, and pay high wages; (ii) "Red firms" hire only red workers, offer safe careers with seniority-based promotions, are small, and pay low wages. A large bias against red workers is socially inefficient and harms both worker types. Although firms individually prefer to be unbiased, paradoxically, they collectively benefit from the bias, as *equilibrium profits increase with the bias*. Our model provides numerous testable predictions, which can guide future empirical work on promotions, discrimination, and inequality.

Keywords: Promotions, Discrimination, Inequality

*We thank Ashwini Agrawal, David Webb, and seminar participants at LSE and QMUL for comments. Ferreira: London School of Economics, CEPR and ECGI; d.ferreira@lse.ac.uk. Nikolowa: Queen Mary University of London; r.nikolowa@qmul.ac.uk. Pikulina: elena.pikulina@sauder.ubc.ca, Finance Division, Sauder School of Business, University of British Columbia.

1 Introduction

There are significant gender promotion gaps in high-skilled occupations. Promotion gaps have been documented for central bankers (Hospido, Laeven, and Lamo (2019)), academic economists (Bosquet, Combes, and García-Peñalosa (2019)), managers of large publicly-listed firms (Cziraki and Robertson (2021)), lawyers (Azmat, Cuñat, and Henry (2024)), and bankers (Bircan, Friebel, and Stahl (2023); Huang, Mayer, and Miller (2023); Ceccarelli, Herpfer, and Ongena (2024)), among others. Much of the empirical literature attributes these gaps to biased promotion decisions. For example, Benson, Li, and Shue (2024) find that women in management-track careers have better pre-promotion performance than men, but men are promoted more often based on subjective assessments of greater "potential." Pikulina and Ferreira (2024) argue that managerial discretion in subjective assessments makes discrimination in promotions hard to detect.[1] They call biased decisions based on subjective assessments *subtle discrimination*.

The evidence on biased promotions raises several questions: Does discrimination in promotions cause inequalities in career opportunities, wages, and welfare? Does it affect labor demand, firm size, and economic efficiency? Does it benefit or harm firms? Addressing these questions requires a theoretical framework to understand the labor market consequences of biased promotions.

We present a model of promotions where vacancies, wages, and internal career paths are endogenously determined in a labor market equilibrium. In the model, homogeneous firms compete for equally productive workers. Workers differ only by a payoff-irrelevant label: "blue" or "red." Firms offer long-term employment contracts, possibly with embedded promotion opportunities. Crucially, when selecting among similarly qualified promotion candidates, firms are biased towards blue workers.

In equilibrium, firms offer contracts in which wages increase upon promotion. However, unlike traditional tournament models, our model has no informational frictions.

[1] Hoffman et al. (2018) show that managerial discretion when selecting candidates reflects mostly biases instead of information.

Thus, the promotion premium is not explained by incentive provision (as in Lazear and Rosen (1981)) or talent retention (as in Waldman (1984)).

Our model follows Ferreira and Nikolowa (2024), who propose a new explanation for the prevalence of employment contracts with promotion opportunities. The model features homogeneous firms and workers, perfectly competitive labor markets with no matching frictions, and rational beliefs. Following Rosen (1986), they assume that workers value the nonpecuniary attributes of their jobs. Firms are endowed with a limited number of high-quality jobs. Because jobs are indivisible and there are more workers than high-quality jobs, workers must be assigned to jobs by some randomization mechanism. In this paper, we call this randomization mechanism a *promotion lottery*. Firms design promotion lotteries by creating several entry-level jobs, which are assigned to young workers. After working in such jobs, a limited number of workers are promoted to the top job. In the absence of frictions, firms randomly choose which workers to promote. In equilibrium, firms create an efficient number of jobs, and all workers have the same ex-ante utility.

The first friction we introduce is an inability to commit to a fair promotion lottery. Since firms are biased in favor of blue workers, this lack of commitment implies that red workers are handicapped when competing with blue workers. We show that, in equilibrium, firms become segregated and hire only blue or red agents. However, to sustain this equilibrium, firms must offer overtly discriminatory contracts: they must write contracts for a given worker type. While such an equilibrium is efficient and egalitarian, both segregation and discriminatory contracts are likely to be socially unacceptable for reasons (such as externalities) that we intentionally left outside the model.

In our main model, we assume – realistically – that overtly discriminatory contracts are infeasible, perhaps because they are illegal. The equilibrium displays a large number of predictions with empirical content:

1. Two types of firms emerge in equilibrium: "mixed firms," which hire both blue and red workers, and "red firms," which hire only red workers.

2. Mixed firms offer risky promotion opportunities, i.e., only a few of the entry-level workers will be promoted (at the firm's discretion) to the next level. By contrast, red firms offer "safe career paths," where workers are promoted based on bureaucratic rules (e.g., seniority).

3. Mixed firms pay high wages, and red firms pay low wages.

4. Mixed firms offer higher promotion premiums than red firms.

5. Mixed firms are larger (i.e, they hire more people) than red firms.

6. Blue workers constitute the majority in mixed firms.

7. In mixed firms, blue workers have a higher probability of promotion than red workers, i.e., there is a promotion gap.

We also consider how an increase in the bias affects equilibrium quantities. If firms become more biased, wages increase in mixed firms and decrease in red firms. Mixed firms also become larger and hire more blue workers, while the total number of mixed firms falls.

We show that the favored group (blue workers) may, paradoxically, be harmed by a bias in its favor. A larger bias weakens the workers' bargaining position, thus giving firms the upper hand in the labor market. Consistent with this result, *profits increase with the bias*. Intuitively, the bias works as an implicit collusion device in the labor market by depressing red workers' outside options, allowing firms to capture a larger share of the surplus. As an unexpected result, the model can rationalize postmodern or Neo-Marxist theories that view discrimination in the labor market as a way of preventing workers from uniting against capitalist exploitation. Ironically, we obtain this result in a model of perfect competition under rational expectations.

2 Setup and Benchmark

In this section, we present the model setup and the benchmark results when firms are unbiased.

2.1 Firms, Workers, and Contracts

We consider an economy with a mass F of identical infinitely-lived firms. Each unique firm is indexed by $\tau \in [0, F]$. A firm can employ workers to perform two tasks, l and h. Tasks are indivisible, and an agent can perform only one of them during each period. We call such tasks *jobs*. At each period, a firm must employ a unit mass of workers in job h. Firms can also employ a mass $n \geq 1$ of workers in job l. There are no slot constraints for job l; firms can create as many l jobs as they wish, i.e., n is a choice variable. Firms cannot change the number of h jobs, which is fixed at 1.

At each period, a mass $E > F$ of young agents enters the labor force and retires after two periods. Each agent has a label, $i \in \{b, r\}$, for blue and red. Labels are observable and productively irrelevant. The proportion of type-i agents in the population is α_i, with $\alpha_b + \alpha_r = 1$. Firms attract job applicants by advertising (and committing to) a number of vacancies for job l (the entry-level job), (n_b, n_r), where n_i is the number of vacancies for type-i workers, and a *career wage schedule*, $(w_{lb}, w_{lr}, w_{hb}, w_{hr})$, where w_{ji} is the (per unit mass of workers) wage for job j if performed by a worker of type i. The total number of vacancies is $n = n_b + n_r$. Firms design internal careers such that all workers start at job l when young. After one period, a unit mass of these workers is promoted to job h, and $n - 1$ workers are fired. Firms retain discretion over promotion decisions.

With a slight abuse of terminology, we call $c = (c_b, c_r)$ an *employment contract*, where $c_i = (n_i, w_{li}, w_{hi})$ for $i \in \{b, r\}$. Because such up-or-out contracts are independent of ability and past performance, we also call them *promotion lotteries*. If $n = 1$, the promotion lottery is degenerate; all entry-level workers are promoted to job h after one period with probability 1. Thus, workers expect a *safe career path*. If $n > 1$, not all workers can be

promoted. Let $p_i(c)$ be the belief that a type-i worker has about his/her promotion probability under contract c. If $p_i(c) < 1$, a type-i worker faces a *risky career path*. Belief $p_i(c)$ is an endogenous function to be determined in equilibrium.

A firm receives revenue $R_j > 0$ for each unit mass of workers employed in job j. A firm's *per contract* profit (assuming no discounting) is thus

$$\pi(c) = R_h - w_h(c) + nR_l - n_b w_{lb} - n_r w_{lr}, \tag{1}$$

where $w_h(c)$ is the total wage bill for the top job. Because the firm chooses who to promote, it is rational to promote first those with lower cost, so we have

$$w_h(c) = \begin{cases} \min\{n_b, 1\} w_{hb} + \max\{1 - n_b, 0\} w_{hr} & \text{if } \ w_{hb} \leq w_{hr} \\ \min\{n_r, 1\} w_{hr} + \max\{1 - n_r, 0\} w_{hb} & \text{if } \ w_{hb} > w_{hr} \end{cases}. \tag{2}$$

Note that when $w_{hb} = w_{hr}$, the firm is indifferent as to whom to promote, and thus, any promotion rule is rational.

Agents derive utility from money and the attributes (i.e., amenities) of their jobs. Let $u_j(w)$ denote the per period utility of having job j and earning wage w, with $u_j'(w) > 0$ and $u_j''(w) < 0$. All agents – regardless of their labels – have the same preferences. We assume $u_h(w) > u_l(w)$ for all w. That is, all else constant, workers prefer job h to l. This preference is due to differences in job characteristics such as prestige, flexibility, autonomy, working conditions, location, or any other nonpecuniary attribute. For simplicity only, from now on, we set $u_l(w) = u(w)$ and $u_h(w) = \theta u(w)$, where $\theta > 1$ and $u'(0) = \infty$.[2] We denote the (per period) utility of an unemployed agent by $\underline{u}$. We assume

Assumption 1. *(i)* $u(R_l) + \theta u(R_h) > 2\underline{u}$*, and (ii)* $u(R_l) < \underline{u}$.

Assumption 1(i) is necessary for promotion lotteries to be efficient. Assumption 1(ii) implies that firms do not want to offer l jobs as "standalone jobs;" l jobs exist only as part

[2]Our results also go through under more general utility specifications $u(\theta, w)$, provided that $u_\theta > 0, u_w > 0, u_{ww} < 0$ and $u_{\theta w} > 0$ (see Ferreira and Nikolowa (2024)).

of a promotion lottery leading to job h. Without loss of generality, to save on notation, from now on we normalize the unemployment utility to zero: $\underline{u} = 0$.

We say that a contract is *feasible* for type i if $n_i > 0$. We define a worker's lifetime utility under a feasible contract c as

$$U_i(c) = u(w_{li}) + p_i(c)\theta u(w_{hi}). \tag{3}$$

Type-i workers expect to be promoted with probability $p_i(c)$ and to be fired when old with probability $1 - p_i(c)$, in which case they become unemployed and enjoy utility $\underline{u} = 0$. We assume that workers have rational expectations: $p_i(c)$ is the actual probability of being promoted given contract c. Note that we impose that beliefs depend only on contract c, and not on other contracts $c' \neq c$ that might be available but were not chosen by the worker. In addition, we assume that all workers and firms share the same beliefs.

Comments on the assumptions. Why would firms offer promotion lotteries? Ferreira and Nikolowa (2024) show that, under similar assumptions on preferences and technology, promotion lottery contracts are strictly optimal for sufficiently high θ. Thus, assuming such contracts is without loss of generality under the implicit assumption that θ is (sufficiently) high. The optimality of promotion lotteries follows from the assumptions that jobs are indivisible and utilities depend on job assignment. Crucially important is that job assignments affect the marginal utility of income. The optimality of lotteries for assigning bundles of money and indivisible goods to people has also been shown in other contexts, such as occupational choice (Bergstrom (1986)), investments in education (Freeman (1996)), and location decisions (Rosen (1997)).

To streamline the analysis, we have assumed that a dismissed worker receives no compensation from the firm. While severance compensation can sometimes be optimal (as Ferreira and Nikolowa (2024) show), its availability does not affect the optimality of promotion lotteries or the qualitative features of the equilibrium. Similarly, the "out" part of the up-or-out contract is not strictly necessary and should not be taken literally. It is

straightforward to extend the model to allow unpromoted workers to remain at the firm earning a wage lower than w_{hi}. Such an extension has no consequences for the model predictions.

2.2 Matching Firms and Workers

We model the matching process as follows. First, all firms simultaneously post (and commit to) contracts. Each firm can post a single contract. Then, all young workers simultaneously apply to as many firms as they wish. Let A_i be the set of the best contracts for type i, i.e., all contracts that give type i the highest utility. The matching mechanism works in two steps. First, the mechanism allocates all contracts in A_i but not in $A_b \cap A_r$ to workers of type $i \in \{b, r\}$. If there are more workers than vacancies, the mechanism selects workers randomly (the same happens if there are more vacancies than workers). In the second step, the mechanism randomly allocates all contracts in $A_b \cap A_r$ to the remaining workers, independently of their types. All remaining unmatched workers become unemployed.

2.3 Equilibrium Conditions

Firms and workers are infinitesimal and, thus, do not behave strategically. Let $c_\tau = (n_{b\tau}, w_{lb\tau}, w_{hb\tau}, n_{r\tau}, w_{lr\tau}, w_{hr\tau})$ denote the contract offered by firm τ. Let $\mathcal{C}$ denote the space of all contracts, $C \subset \mathcal{C}$ a set of contracts c_τ for all $\tau \in [0, F]$, and $C_i \subseteq C$ the set of all feasible contracts in C for $i \in \{b, r\}$. Define

$$U_i^*(C_i) = \max_{c \in C_i} U_i(c). \tag{4}$$

That is, $U_i^*(C_i)$ is the maximum utility a type-i worker obtains when free to choose any contract in C_i. In equilibrium, $U_i^*(C_i)$ must equal a type-i's worker reservation utility; no type-i worker would apply to contract c if $U_i(c) < U_i^*(C_i)$

We say that the labor market is *tight* if all agents are employed. In a similar setup, Fer-

reira and Nikolowa (2024) show conditions under which a tight-labor-market equilibrium exists. In this paper, we assume that these conditions are met. Thus, we restrict attention to tight labor markets.

Definition 1. An equilibrium is a set of contracts $C^* = \{c_\tau^* : \tau \in [0, F]\}$ such that

1. If $c_\tau^* \in C^*$, then $c_\tau^* \in \arg\max_{c \in \mathcal{C}} \pi(c)$ subject to $U_i(c) \geq U_i^*(C_i^*)$, for $i \in \{b, r\}$;
2. $U_i^*(C_i^*) \geq 0$;
3. $\int_0^F n_{i\tau}^* d\tau = \alpha_i E$ for $i \in \{b, r\}$.

Condition 1 says that firms choose contracts to maximize profits, while taking all other firms' contracts as given. Condition 2 implies that young workers accept the equilibrium contracts if and only if these workers are better off employed. Condition 3 requires labor demand to equal labor supply.

2.4 Benchmark: Unbiased firms

As a benchmark, we assume that there is only one type of worker, say b, so that $\alpha_b = 1$; we drop subscript i for notational simplicity. If a firm offers contract $c = (n, w_l, w_h)$, all agents must have the same probability of promotion $p(c) = \frac{1}{n}$. A worker's lifetime utility is given by

$$U(c) = u(w_l) + \frac{1}{n}\theta u(w_h), \tag{5}$$

and a firm's per contract profit is

$$\pi(c) = R_h - w_h + n(R_l - w_l). \tag{6}$$

Note that R_h is simply a profit shifter, which has no impact on the equilibrium contracts. Thus, R_h is a "free parameter," which we assume to be sufficiently high so that Assumption 1 holds and firms are willing to operate (i.e., profits are non-negative). Let Θ denote

the set of parameters (R_l, θ, E, F) such that a tight-labor-market equilibrium exists.[3] The next proposition shows that a tight-labor-market equilibrium is always unique.

Proposition 1. *If $(R_l, \theta, E, F) \in \Theta$, there is a unique equilibrium contract $c^{FB} = (n^*, w_l^*, w_h^*)$ offered by all firms.*

The equilibrium is determined by three intuitive conditions (see the proof of Proposition 1 for the calculations). First, the marginal utilities in each job must equalize:

$$u'(w_l^*) = \theta u'(w_h^*). \tag{7}$$

This condition implies that *wages must increase upon promotion*. A promotion premium exists because the marginal utility of money is higher at the top job. Promotion tournament models also predict the existence of a promotion premium, but for different reasons. In classic tournament models in the tradition of Lazear and Rosen (1981), promotion premiums exist to provide workers with incentives to exert non-contractible effort. In market-based tournaments (Waldman (2013)), promotion premiums exist to retain promoted high-skilled workers. By contrast, in promotion-lottery models – as in this paper – workers are (productively) identical, and there is no non-contractible action. Promotion premiums exist because cost-minimizing contracts must equalize workers' marginal utilities across jobs.[4]

Second, we have

$$R_l = w_l^* + \frac{1}{n^*} \frac{\theta u(w_h^*)}{u'(w_l^*)}. \tag{8}$$

This condition equates the marginal productivity of job l, R_l, to the marginal cost of hiring

[3] Ferreira and Nikolowa (2024) show that Θ is non-empty. More generally, an equilibrium always exists for any set of parameters, but it may involve some unemployed workers or idle firms.

[4] At first blush, condition (7) may appear obvious. However, an individual worker would prefer instead the equalization of expected marginal utilities: $u'(w_l) = p\theta u'(w_h)$. That is, if (7) holds, a worker would be willing to accept a small reduction in w_h^* in exchange for a similar increase in w_l^*. We obtain (7) because if the firm increases w_l^* by ε, it costs $n\varepsilon = \frac{\varepsilon}{p}$, and the firm must reduce w_h^* by this amount to break even. Thus, the relevant per worker condition is $u'(w_l)\varepsilon = p\theta u'(w_h)\frac{\varepsilon}{p}$, which implies (7).

young workers. An additional worker costs the firm w_l^* plus a "share" $1/n^*$ of the utility in job h, $\theta u(w_h^*)$, which is converted into dollars by dividing it by the marginal utility of money, $u'(w_l^*)$. Intuitively, if (8) does not hold, firms would like to increase or reduce the number of job positions they offer. Condition (8) is analogous to the requirement for equalizing wages and marginal productivities in spot labor markets. It implies that, under a promotion-lottery contract, the entry-level wage must be lower than the marginal productivity of young workers.

Finally, supply and demand imply that $n^* = E/F$. Proposition 1 implies a single solution (n^*, w_l^*, w_h^*). Figure 1 shows (w_l^*, w_h^*) as a function of R_l for two values of θ, for utility $u(w) = \ln w$ and $E = 2F$ (which implies $n^* = 2$). For visualization, we use different scales for w_l^* and w_h^*.

Figure 1 shows that both wages increase with R_l – the marginal productivity of job l – but also that the top wage is much more sensitive to R_l than the entry-level wage (please note the different scales). This fits with anecdotal accounts of early-career workers in professional services performing menial tasks for wages that are "insensitive" to the actual marginal productivity of such tasks. As a consequence, an increase in R_l increases *within-firm* wage inequality, which is an important empirical prediction of the model.

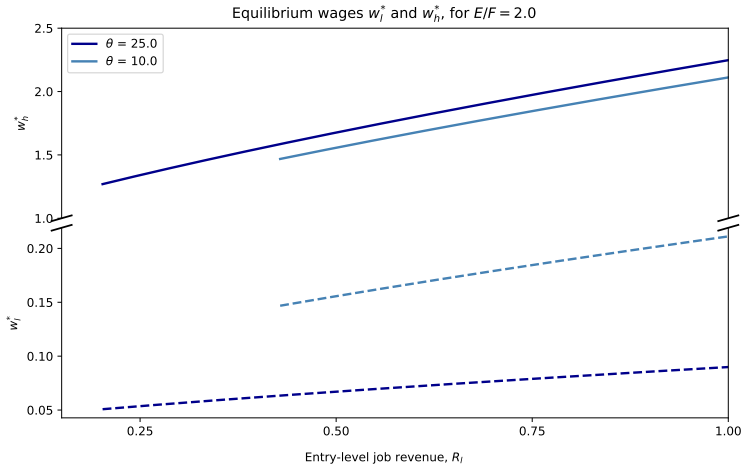


Figure 1: Equilibrium wages in the benchmark case.

Note: This figure presents the benchmark equilibrium wages w_h^* (solid lines) and w_l^* (dashed lines), for top job h and entry-level jobs l, as functions of the entry-level job productivity R_l, for $E/F = 2.0$

An increase in θ (the "prestige" of the top job) also widens the wage gap between entry-level and top jobs. Thus, more prestigious high-level jobs increase within-firm wage inequality, which is another unique prediction of the model.

Overall, Figure 1 illustrates that promotion lotteries produce significant ex-post inequality among ex-ante identical workers. But this is "good" inequality: promotion lotteries represent a Pareto improvement upon deterministic employment contracts.[5] To see that the equilibrium is efficient, notice first that Assumption 1 implies that combining jobs l and h in one career is efficient, while separating them is not. Second, because the supply of h jobs is $F < E$, and jobs are indivisible, the assignment of workers to job h must involve some kind of lottery. Finally, efficiency requires equalizing marginal utilities as in (7) and equalizing marginal productivity to marginal cost as in (8). An equilibrium with perfectly competitive firms that offer promotion lotteries is efficient because these conditions are met.

Suppose now that $\alpha_b < 1$. Since blue and red workers are productively identical, the

[5]Rosen (1997) calls this type of efficient lottery-driven inequality "manufactured inequality."

equilibrium described in this section remains an equilibrium. However, because firms can now offer different contracts to different types, other equilibria may also exist. In particular, if $w_{hb} = w_{hr}$, the firm is indifferent between which type to promote. If the firm is unbiased (i.e., $p_b = p_r = 1/n$), the equilibrium contract in Proposition 1 is the only one feasible. This contract is offered to all agents irrespective of type, and the equilibrium distribution of types across firms is indeterminate. However, other equilibria may arise if firms are *biased*. We consider this possibility in the next section.

3 Biased Firms

This section considers the case in which firms can be biased towards one type of worker. Following Pikulina and Ferreira (2024), we say that a firm is *subtly biased* towards blue workers if it is more likely to promote the blue worker when indifferent between blue and red workers. Notice that any promotion decision is rational when the firm is indifferent between blue and red workers. A biased promotion lottery is a form of *subtle discrimination*: an act that cannot be objectively ascertained as discriminatory. The idea is that subjective assessments make discrimination in promotions hard to detect.

3.1 Biased Promotion Probabilities

Pikulina and Ferreira (2024) consider the case of two identical workers, Blue and Red, who compete for a promotion. The firm is biased towards Blue: the promotion probabilities are $p_b = \frac{1+\beta_b}{2}$ and $p_r = \frac{1-\beta_b}{2}$ for Blue and Red, respectively, where $\beta_b \in [0, 1]$. Let $p_b - p_r = \beta_b \geq 0$ denote the *subtle bias* towards Blue. Pikulina and Ferreira (2024) rationalize β_b as the limiting case of a model where workers have small observable differences and a biased decision-maker has private information about the workers' productive abilities. The bias is subtle because, given the near-identical observable worker productivities, the decision-maker may use his private information to justify any promotion decision. In

other words, because promotion decisions are partly based on subjective assessments, the decision-maker can plausibly deny being biased. A subtle bias implies an inability to commit to a particular promotion probability. We can think of subtle discrimination as a friction arising from this inability to commit.

Here, we extend Pikulina and Ferreira's notion of subtle discrimination to the case of n agents. Suppose the firm is biased towards blue workers and let q denote the proportion of blue workers competing for the top job. Let $p_i(n,q)$ denote the probability that a given type-i agent is promoted, $i \in \{b,r\}$. The probability that the promoted agent is of type i is thus $n_i p_i(n,q)$, where n_i is the number of type-i workers competing for promotion. We define the bias in the general (n,q) case as:

$$\beta_b(n,q) = n\left[qp_b(n,q) - (1-q)p_r(n,q)\right] - (2q-1). \tag{9}$$

The first term on the right-hand side is the probability that the promoted worker is blue minus the probability that the promoted worker is red. The second term is the same difference of probabilities if the firm is unbiased. Note that if $p_b(2,0.5) = \frac{1+\beta_b}{2}$ and $p_r(2,0.5) = \frac{1-\beta_b}{2}$, we have $\beta_b(2,0.5) = \beta_b$, which is the bias in the two-worker, two-type case.

To pin down $p_i(n,q)$, we take an axiomatic approach and impose intuitive three properties on this function.

Axiom 1 (Monotonicity)**.** *$p_i(n,q)$ decreases with q.*

Axiom 1 simply says that, all else constant, both blue and red workers find it more difficult to be promoted when there are more blue workers. Any reasonable function $p_i(n,q)$ must have this property.

Axiom 2 (Size-invariance)**.** *$p_i(n,q)$ is homogeneous of degree -1 in n: $p_i(zn,q) = \frac{p_i(n,q)}{z}$ for all $z > 0$.*

In an unbiased promotion lottery, doubling the number of workers cuts everyone's

probability of promotion by half. That is, in unbiased promotions, the promotion probability is homogeneous of degree minus one in the number of candidates. Axiom 2 preserves this property in biased promotion settings, as long as the proportions of each type remain constant. Intuitively, we assume that the firm's bias is invariant to firm size; the firm does not become more or less biased as it grows in size. That is, Axiom 2 implies that $\beta_b(q, n)$ is independent of n.

Axiom 3 (Constant returns)**.** *Probability changes due to proportion changes are independent of proportion levels: $p_i(n, q_0 + \epsilon) - p_i(n, q_0)$ is independent of q_0 for any ϵ.*

Intuitively, Axiom 3 implies that the firm does not become intrinsically more or less biased as the proportion of blue workers changes. Changes in q affect the promotion probabilities only due to the need to compete against more blue workers. The following proposition shows that Axioms 1-3 are sufficient to determine function $p_i(n, q)$.

Proposition 2. *Axioms 1-3 imply that the individual promotion probabilities are*

$$p_b(n, q) = \frac{1 + (1 - q)\beta}{n} \text{ and } p_r(n, q) = \frac{1 - q\beta}{n}, \tag{10}$$

where $\beta = 2\beta_b$, and $\beta_b \in [0, 1]$ is the bias in the $(2, \frac{1}{2})$ case.

Although promotion probabilities are uniquely defined given Axioms 1-3, they are consistent with multiple mechanisms for choosing whom to promote. For example, the following two-step procedure delivers (10): (i) choose which type to promote via a biased binary lottery, with the probability of promoting a blue worker equal to $q(1 + (1 - q)\beta)$; (ii) randomly choose one worker from the group selected in step (i).

3.2 Equilibrium with Overt Discrimination

If firms are biased, workers care about the identity of their co-workers, as they will compete for the same position. Suppose firms hire workers of both types. Then, a firm must

offer the same wage for the high-level job to both types: $w_{hb} = w_{hr}$. If not, the firm will always promote the cheapest type, and the other type will not accept the contract. Under the same top wage, firms now have a preference for promoting blue workers. Thus, under the same contract, workers – both blue and red – prefer to work in firms with fewer blue workers.

An equilibrium where some firms have a mixed workforce, with blue and red workers, does not exist. To see this, note that a firm would want to hire both worker types only if they cost the same. However, because in a mixed firm a red worker has a lower probability of promotion than a blue worker, red workers have lower utility than blue workers. Thus, the firm would strictly prefer to hire only red workers, who would accept a contract at lower career wages, provided they do not have to compete with blue workers for the promotion. We then have the following result.

Proposition 3. *If firms are subtly biased, ($\beta \geq 0$), there is a unique equilibrium where only two contracts are offered, $c' = \left(c^{FB}, (0,0,0)\right)$ and $c'' = \left((0,0,0), c^{FB}\right)$. The equilibrium is symmetric and displays full segregation: $\alpha_b F$ firms hire only blue workers, and $(1-\alpha_b)F$ firms hire only red workers.*

Proposition 3 implies that, as long as firms can overtly discriminate by offering different contracts to blue and red workers, subtle biases do not matter. The unique segregated equilibrium is efficient and egalitarian: both blue and red workers expect the same wages and have the same expected utility. Still, type segregation is an unappealing property. What remains unmodeled are the potential externalities from segregation. For example, segregation may reinforce stereotypes and biases, with negative societal consequences. It may also affect the perceived quality of the different jobs, here assumed to share the same θ. With different job qualities, wages would also differ, and inequalities in career opportunities would arise.

For several good reasons, most advanced economies have made overt discrimination illegal. An unintended casualty of this trend is that firms cannot use employment con-

tracts to "de-bias" their promotion practices.

3.3 Equilibrium When Overt Discrimination Is Not Possible

We now turn to the main version of our model, where we assume – realistically – that firms cannot offer discriminatory employment contracts. Each firm $\tau \in [0, F]$ must offer a non-discriminatory contract, $c_\tau = (n_\tau, w_{l\tau}, w_{h\tau})$, and accept both types of workers. If more than n_τ workers apply to the firm, it must randomly choose among them, regardless of type.

Firms cannot commit to an unbiased promotion selection and are thus subject to subtle biases. The differences in the importance of subjective information for hiring and promotion decisions explain this asymmetry. When hiring for entry-level jobs, firms have little private information about the applicants and are thus likely to choose mainly based on observable characteristics. In contrast, in promotion decisions, firms arguably have substantial subjective information about the candidates. This subjectivity in assessment allows firms to plausibly deny being biased.

Now, when applying for a job, a worker cares not only about the contract but also about the expected type composition of the firm's workforce. Let $q : \mathcal{C} \to [0, 1]$ denote the proportion of blue workers expected to be hired under a contract $c \in \mathcal{C}$, where $\mathcal{C}$ is the space of all feasible contracts. Function $q(.)$ represents agents' *beliefs*. We assume that beliefs $q(.)$ are rational and shared by all workers and firms. Firms cannot choose the shape of $q(.)$; this shape is determined in equilibrium by rational expectations. Note that beliefs depend only on the relevant contract c.

A worker's lifetime utility under c is now

$$U_i(c) = u(w_l) + p_i(n, q(c))\theta u(w_h), \tag{11}$$

where function $p_i(n, q)$ is given by (10). Let $m_{i\tau}(C)$ denote the mass of workers of type i matched with firm τ by the matching mechanism for a given set of contracts C. We modify

the definition of the equilibrium to impose the constraint of no overt discrimination:

Definition 2. An equilibrium without overt discrimination is a set of contracts $C^* = \{c_\tau^* : \tau \in [0, F]\}$ and a belief function $q : \mathcal{C} \to [0, 1]$ such that

1. If $c_\tau^* \in C^*$, then $c_\tau^* \in \arg\max_{c \in \mathcal{C}} \pi(c)$ subject to $U_i(c) \geq U_i^*(C^*)$, for $i \in \{b, r\}$;
2. If firm $\tau \in [0, 1]$ offers $c_\tau \in \mathcal{C}$ such that $U_i(c_\tau) \geq U_i^*(C^*) \geq 0$ for at least one $i \in \{b, r\}$, then $m_{b\tau}(C^* \cup \{c_\tau\}) = n_\tau q(c_\tau)$ and $m_{r\tau}(C^* \cup \{c_\tau\}) = n_\tau(1 - q(c_\tau))$;
3. $\int_0^F n_\tau^* d\tau = E$.

Condition 1 is the standard requirement that firms must maximize profit, taking the workers' equilibrium utilities as given. Condition 2 states that an equilibrium contract must maximize the utility of at least one type of worker. Condition 3 states that labor demand must equal labor supply.

Condition 2 also imposes constraints on the belief function, $q(.)$. If c_τ^* is an equilibrium contract, beliefs must be consistent with the matching mechanism in equilibrium, implying $m_{b\tau}(C^*) = n_\tau^* q(c_\tau^*)$. That is, beliefs must be consistent with equilibrium play. Furthermore, if a single firm deviates and offers $c_\tau \notin C^*$, belief $q(c_\tau)$ must also be consistent with the matching mechanism, assuming all other firms offer their equilibrium contracts. In sum, we require rational expectations both in equilibrium and following a single deviation from the equilibrium.

We can now characterize the equilibrium contracts. We say that an equilibrium $(C^*, q(.))$ displays full segregation if $q(c_\tau^*) \in \{0, 1\}$ for all $\tau \in [0, F]$. That is, in a full segregation equilibrium, there are no mixed firms. The next result shows that there is no full segregation equilibrium when firms cannot offer discriminatory contracts.

Lemma 1. *If overt discrimination is not possible, there is no equilibrium with full segregation.*

Intuitively, segregation cannot happen because blue workers would all apply to firms with only red workers.

Lemma 2. *Suppose* $(C^*, q(.))$ *is an equilibrium.*

1. *There must exist* $c' \in C^*$ *such that* $n' > 1$ *and* $q(c') \in (0, 1)$.

2. *If* $c'' \in C^*$ *and* $c'' \neq c'$, *then* $n'' = 1$.

Lemma 2 implies that there are, at most, two contracts in equilibrium. One contract is a *risky career path*, i.e., a promotion lottery where not all entry-level workers can be promoted. Firms that offer such contracts must be mixed: the proportion of blue workers cannot be exactly zero or one. A *safe career path* contract (where $n = 1$) may also exist. Firms offering such contracts promote all their entry-level workers. The next result shows that safe career path firms hire only red workers.

Lemma 3. *In equilibrium, blue workers must choose a contract where* $n > 1$.

The next proposition is a direct consequence of Lemmas 1-3.

Proposition 4. *Let* $(C^*, q(.))$ *denote an equilibrium. The equilibrium displays* ***partial segregation****:*

1. *A fraction* $s^* \in (0, 1]$ *of "mixed firms" hire* $n^* > 1$ *workers for wages* (w^*_{lm}, w^*_{hm})*; the proportion of blue workers in mixed firms is* $q^* \in [\alpha_b, 1)$.

2. *A fraction* $1 - s^*$ *of "red firms" hire one red worker each for wages* (w^*_{lr}, w^*_{hr}).

This proposition implies that, in equilibrium, firms differentiate themselves by choosing one of two contracts. "Mixed firms" employ young blue and red workers by offering them risky career paths, where only a few workers are promoted to the top job when they are old. "Red firms" employ only young red workers, who are then promoted with probability one to the next level. These safe career paths can be interpreted as bureaucratic promotion rules, such as promotion by seniority, which are common in some government and nonprofit organizations. An equilibrium must have mixed firms but may or may not have red firms.

Proposition 4 implies that an equilibrium is a set of two unique contracts $\{c_r^*, c_m^*\}$ and a belief function $q(c)$, where firm $\tau_m \in [0, s^*F]$ offers contract $c_m^* = (n_m^*, w_{lm}^*, w_{hm}^*)$, firm $\tau_r \in (s^*F, F]$ offers contract $c_r^* = (1, w_{lr}^*, w_{hr}^*)$, and $s^* \in (0, 1]$. The equilibrium must satisfy seven conditions:

(i) Workers' lifetime utilities are

$$\begin{aligned} U_r^* &:= u(w_{lm}^*) + \tfrac{1-q^*\beta}{n_m^*}\theta u(w_{hm}^*) \\ U_b^* &:= u(w_{lm}^*) + \tfrac{1+(1-q^*)\beta}{n_m^*}\theta u(w_{hm}^*) \end{aligned} \tag{12}$$

where $q^* := q(c_m^*)$.

(ii) Type-m firms choose c_m^* optimally:

$$\pi_m^* = \max_{c_m \in \mathcal{C}} R_h - w_{hm} + n_m(R_l - w_{lm}) \tag{13}$$

subject to

$$u(w_{lm}) + \frac{1-q^*\beta}{n_m}\theta u(w_{hm}) \geq U_r^* \tag{14}$$

$$u(w_{lm}) + \frac{1+(1-q^*)\beta}{n_m}\theta u(w_{hm}) \geq U_b^*. \tag{15}$$

(iii) Type-r firms choose c_r^* optimally:

$$\pi_r^* := \max_{c_r \in \mathcal{C}} R_h - w_{hr} + R_l - w_{lr} \tag{16}$$

subject to

$$u(w_{lr}) + \theta u(w_{hr}) \geq U_r^*. \tag{17}$$

(iv) Firms' choice of contract type must be optimal:

$$\begin{cases} \text{If } \pi_m^* > \pi_r^*, \text{ then } s^* = 1 \\ \text{If } \pi_m^* = \pi_r^*, \text{ then } s^* < 1 \end{cases} \tag{18}$$

(v) There is no profitable deviation: $\nexists\, c^d$ such that $m(C^* \cup \{c^d\}) = n_d q(c^d)$ and $\pi(c^d) > \pi^*_m$.

(vi) If $s^* < 1$, red workers are indifferent between contracts c^*_m and c^*_r:

$$U^*_r = u(w^*_{lr}) + \theta u(w^*_{hr}). \tag{19}$$

(vii) Labor markets clear:

$$\begin{cases} s^* q^* n^*_m F = \alpha_b E \\ (1 - s^* + s^*(1 - q^*) n^*_m) F = (1 - \alpha_b) E \end{cases}. \tag{20}$$

The next result simplifies the calculation of the equilibrium by showing that only the red workers' participation constraint binds in equilibrium.

Lemma 4. *In equilibrium, the red workers' participation constraint (14) binds, and the blue workers' participation constraint (15) does not bind.*

The next result shows that the absence of a profitable deviation can be inferred by checking a single condition.

Lemma 5. *A belief function $q(.)$ such that equilibrium condition (v) holds exists if and only if*

$$\pi^*_m \geq \max_{c \in \mathcal{C}} \pi(c) \;\; s.t. \; U_b(c; 1) \geq U^*_b, \tag{21}$$

where $U_i(c; q)$ is i's utility from contract c and belief q.

3.4 Equilibrium Predictions

As in the benchmark case, the equilibrium is characterized by a few intuitive conditions. First, for mixed firms we have:

$$u'(w^*_{lm}) = \theta(1 - q^* \beta) u'(w^*_{hm}). \tag{22}$$

Note that, as before, wages increase upon promotion. However, unlike the benchmark case, *marginal utilities are not equalized across jobs*. The marginal rate of substitution between wages in each job (i.e., the ratio of marginal utilities) is now $1 - q^*\beta$. Intuitively, the bias β reduces the value of the wage in the top job for red workers, because these workers are less likely to be promoted. Mixed firms choose their optimal contracts hoping to attract red workers, as these are cheaper than blue workers, implying that firms equalize red workers' marginal utilities *as perceived by them*. Because firms cannot overtly discriminate by turning down blue candidates, in equilibrium $q^* n_m^*$ blue workers will show up at each mixed firm's door, and these firms will prefer hiring them to keeping slots unfilled.

Condition (22) implies that the equilibrium is *inefficient*. To achieve the first best, we need $\beta = 0$, i.e., no subtle bias. Thus, we can think of β as a friction: firms with positive subtle biases are unable to commit to a fair promotion lottery. All else constant, this inefficiency manifests itself as a smaller promotion premium than the first best. In other words, the bias against red workers reduces wage inequality between young and old workers within the same firm. But this is a reduction in "good inequality."

The "all else constant" condition is an important qualifier. An increase in β will not only affect the marginal condition in (22); it will also have non-obvious general equilibrium effects. In particular, an increase in β will reduce the equilibrium utility of red workers, U_r^*. Because red workers' outside opportunities are endogenously determined, higher β implies lower outside utilities, which translates into lower bargaining power for these workers. Facing less worker bargaining power, firms can offer cheaper contracts: they can offer lower wages (w_{lm} and w_{hm}) or hire more workers, i.e., they can increase n_m^*, reducing workers' promotion probabilities. Thus, unlike models of spot labor markets, in promotion lottery models, the relationship between U_r^* and wages is, in principle, ambiguous.

The following condition establishes the link between entry-level wages and the red

workers' outside utility:

$$R_l = w^*_{lm} + \frac{U^*_r - u(w^*_{lm})}{u'(w^*_{lm})}. \tag{23}$$

This condition follows from (8) and, thus, it implies that the marginal productivity of workers equals their marginal cost to the firm. Note that the bias does not *directly* affect this condition. But it may affect it indirectly through U^*_r. Because equilibrium existence implies a unique solution to (23), w^*_{lm} decreases with U^*_r.

To illustrate the total effect of changes in β on the equilibrium wages, in Figure 2 we show wages w^*_{lm} and w^*_{hm} using the same parametrization as in Figure 1. We focus on the case in which the equilibrium is interior, i.e., $s^* < 1$. In this case, both types of firms – mixed firms and red firms – exist in equilibrium.[6] Note that the entry-level wage *increases* with the bias. This apparent counterintuitive effect is a consequence of biased promotion lotteries. As firms become more biased against red workers, the equilibrium utility U^*_r falls, which, from (23), implies that entry-level wages increase.

[6]The case where $s^* = 1$ is a corner solution in which all firms are mixed. This equilibrium arises when β – the subtle bias – is sufficiently small. Empirically, this equilibrium is quite similar to the benchmark case with $\beta = 0$, although it also creates substantial inequality between blue and red agents. We offer a full analysis of this equilibrium in the Internet Appendix.

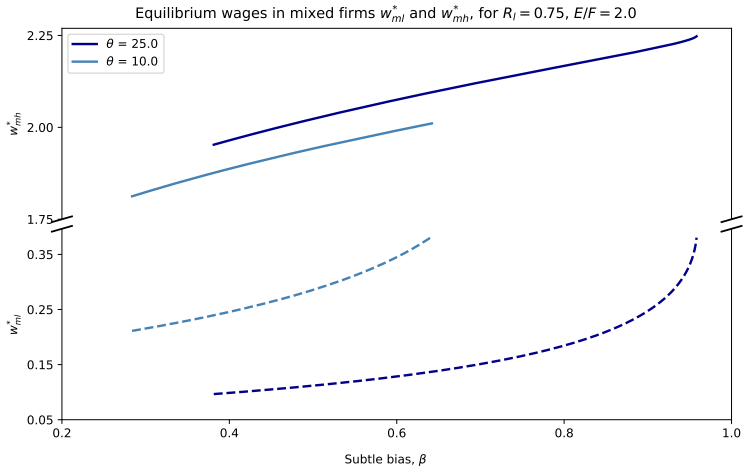


Figure 2: Wages in mixed firms, w^*_{hm} and w^*_{lm}, under subtle bias.

Note: This figure presents the equilibrium wages in mixed firms w^*_{mh} (solid lines) and w^*_{ml} (dashed lines), for top job h and entry-level jobs l, as functions of subtle bias, β, for $R_l = 0.75$, $E/F = 2.0$.

Figure 2 shows that the top wage also increases with the bias. Since both wages increase when workers become less powerful, it must be the case that firms hire more workers, which implies a lower promotion probability. Figure 3 shows that this is indeed the case: Firm size, or the *span of control*, n^*_m, increases with the bias. In other words, as workers become less powerful, firms increase their profits by hiring more workers, lowering promotion probabilities.

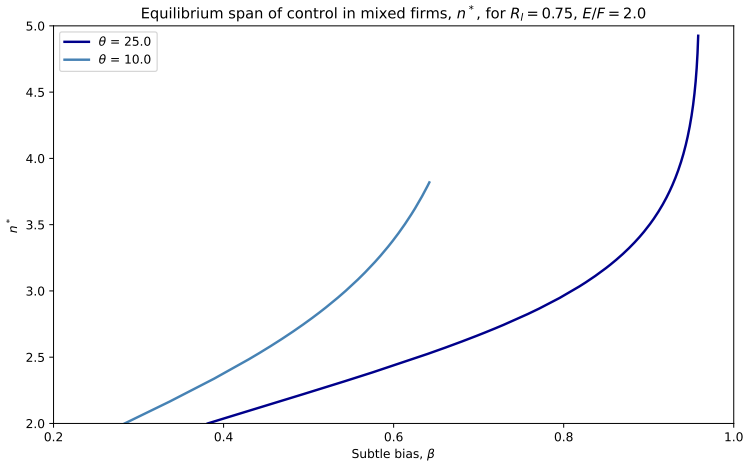


Figure 3: Span of control in mixed firms, n_m^*.

Note: This figure presents the equilibrium span of control in mixed firms, n^*, as a function of subtle bias, β, for $R_l = 0.75$, $E/F = 2.0$.

Under an equilibrium with two types of firms, the within-firm wage gap in mixed firms has an inverted-U shape in the bias, as Figure 4 shows. As β increases, initially, the span-of-control effect dominates, and the top wage increases faster than the entry-level wage, increasing the wage gap. But, as β becomes large, the fraction of blue workers in mixed firms, q^*, also increases (see Figure 5). Note that blue workers are the majority group in mixed firms. As q^* increases, the direct effect of β on w_{mh}^* becomes stronger, which moderates the increase in the top wage, eventually reducing the wage gap.

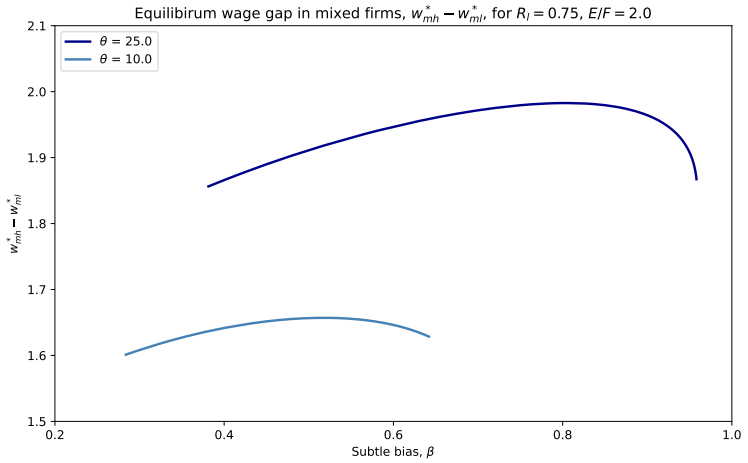


Figure 4: Wage gap between top job h and entry-level jobs l in mixed firms, $w^*_{hm} - w^*_{lm}$.

Note: This figure presents the equilibrium wage gap in mixed firms $w^*_{mh} - w^*_{ml}$, as a function of subtle bias, β, for $R_l = 0.75$, $E/F = 2.0$.

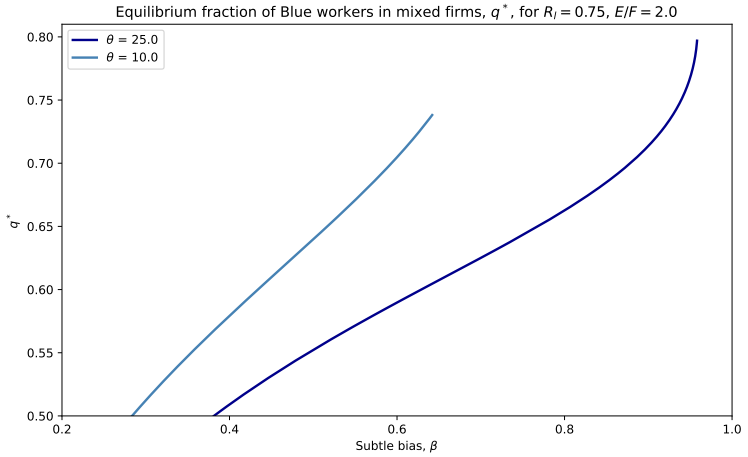


Figure 5: Fraction of Blue workers in mixed firms, q^*.

Note: This figure presents the equilibrium fraction of Blue workers in mixed firms, q^*, as a function of subtle bias, β, for $R_l = 0.75$, $E/F = 2.0$.

We now turn to the other type of firm: red firms. Red firms offer safe career paths to their employees, i.e., the only criterion for promotion is seniority. Red firms cannot

discriminate in hiring; they end up with a fully segregated workforce because only red workers find it optimal to apply to these firms. In equilibrium, red firms also do not discriminate in promotions, implying their career wages equalize red workers' marginal utilities:

$$u'(w^*_{lr}) = \theta u'(w^*_{hr}). \tag{24}$$

Condition (24) implies that wages increase upon promotion in safe career paths, too. Figure 6 shows the career wages in red firms. Note that, at each job level, *red firms pay lower wages than mixed firms*. Thus, the model endogenously implies the existence of high-wage firms and low-wage firms. Note that high-wage firms are larger (i.e., they employ more workers) than low-wage firms. Unlike the case of mixed firms, in red firms both the top and bottom wages decrease with β. In a safe career path contract, the span of control is fixed at 1, implying that a reduction in workers' bargaining power must decrease wages.

The model implies that discrimination in promotions in mixed firms has negative spillovers on workers in non-discriminating firms. These are general equilibrium effects. Subtle biases in promotion in larger, mixed firms distort unfavored agents' early career decisions towards smaller firms with lower wages but more predictable promotion prospects. As more red workers compete for jobs in these firms, wages fall.

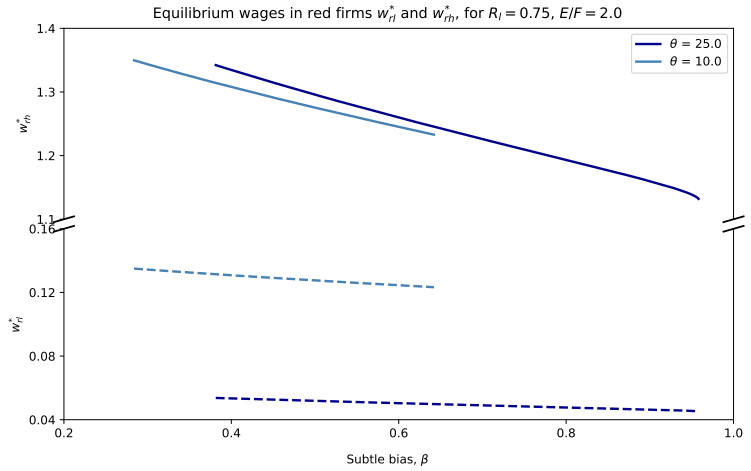


Figure 6: Wages in red firms, w^*_{hr} and w^*_{lr}, under subtle bias.

Note: This figure presents the equilibrium wages in red firms w^*_{rh} (solid lines) and w^*_{rl} (dashed lines), for top job h and entry-level jobs l, as functions of subtle bias, β, for $R_l = 0.75$, $E/F = 2.0$.

While red firms choose wages efficiently given their career structures, the mere existence of these firms indicates resource misallocation because, in the first best, all firms should offer risky career paths. Thus, the existence of red firms reduces "good" inequality. Another way of stating it is that red firms are inefficiently small. By the same token, mixed firms are too large. Intuitively, the existence of subtle biases in the promotion processes of large firms may distort career opportunities in the economy. For higher levels of bias, the number of mixed firms decreases (see Figure 7), while their size increases (see Figure 3). Thus, in equilibrium, we have a small number of inefficiently-large high-wage firms, which hire many favored agents (blue) and a minority of unfavored agents (red). These firms offer risky promotion lotteries leading to very high wages upon promotion. By contrast, smaller, bureaucratic low-wage firms specialize in hiring unfavored workers and offer them predictable careers leading to lower wage increases upon promotion.

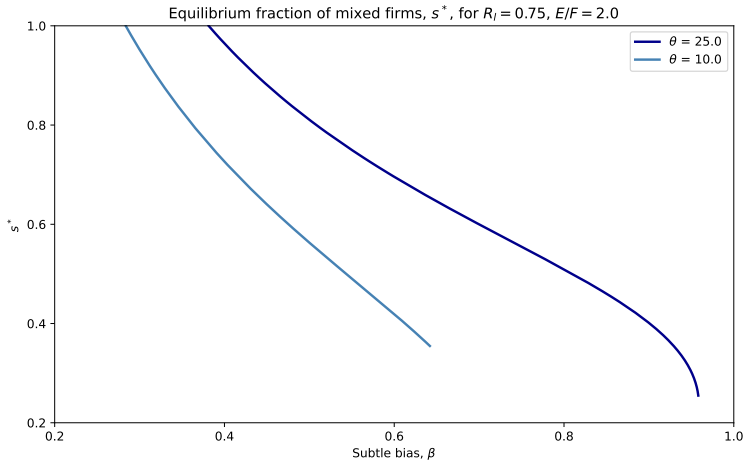


Figure 7: Fraction of mixed firms, s^*.

Note: This figure presents the equilibrium fraction of mixed firms, s^*, as a function of subtle bias, β, for $R_l = 0.75$, $E/F = 2.0$.

Figure 8 shows the average wage for each position (l and h) across all firms in the economy. The figure shows that the average *promotion wage premium* decreases with the bias. Thus, when biases in promotions are large, the age/experience wage gap is lower.

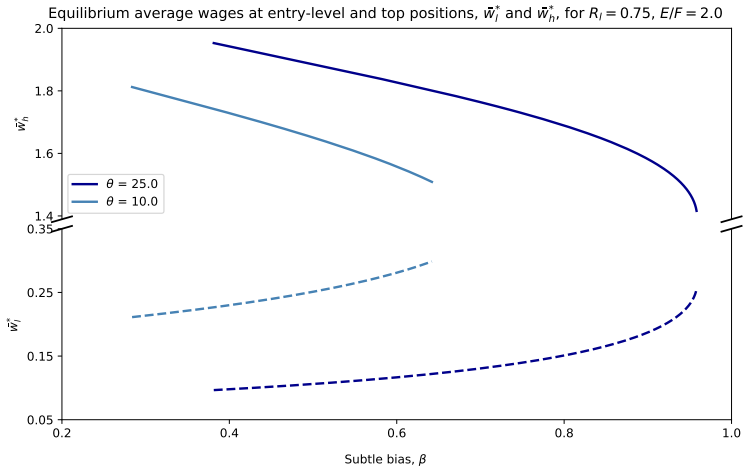


Figure 8: Average wages at entry-level and top positions, $\bar{w}_l^*$ and $\bar{w}_h^*$.

Note: This figure presents the equilibrium average wages at entry-level (dashed lines) and top positions (solid lines) across all firms, $\bar{w}_l^*$ and $\bar{w}_h^*$, as functions of subtle bias, β, for $R_l = 0.75$, $E/F = 2.0$.

The careers in high-wage firms are unambiguously better than those in low-wage firms. Thus, favored workers strictly prefer working for high-wage firms. By contrast, unfavored workers are indifferent between working in large or small firms. Those minority workers who are promoted in high-wage firms will be paid substantially more than those promoted in low-wage firms. Thus, another empirical prediction is that, at any given rank, there is more wage inequality among discriminated than non-discriminated workers.

Since red workers are disproportionally represented in low-wage firms, the average wage of red workers across all firms and positions is lower than that of blue workers (see Figure 9). The blue-red wage gap initially increases with the bias, but then decreases for sufficiently high levels of β.

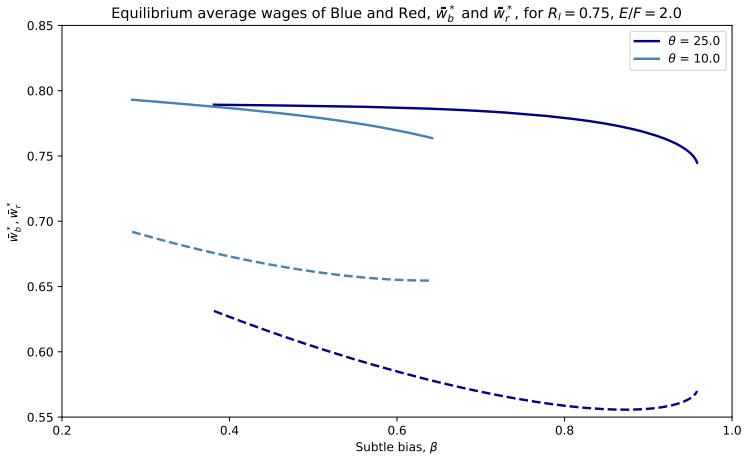


Figure 9: Average wages of Blue and Red workers, $\bar{w}_b^*$ and $\bar{w}_r^*$.

Note: This figure presents the equilibrium average wages of Blue (solid lines) and Red (dashed lines) workers (across all firms and positions), $\bar{w}_b^*$ and $\bar{w}_r^*$, as functions of subtle bias, β, for $R_l = 0.75$, $E/F = 2.0$.

3.5 Welfare

The higher average wage for blue workers is not a compensating differential for their higher promotion risk. Figure 10 shows that blue workers' lifetime utility is higher than that of red workers. Thus, blue workers enjoy rents for having the blue label. Perhaps surprisingly, blue workers' utilities decrease with the bias. As the bias increases, red workers become less powerful. Because mixed firms choose contracts that bind red workers' participation constraints, the employment contracts in such firms become less favorable to all workers when red workers are weaker. Thus, blue workers may also suffer from discrimination in their favor. This competition effect can make blue workers worse off in a biased equilibrium than in the benchmark case of $\beta = 0$. The horizontal lines in Figure 10 shows the utilities of blue and red workers under the no-discrimination benchmark ($\beta = 0$). For sufficiently high β, favored workers are paradoxically worse off in the polarized equilibrium that is biased in their favor than in the egalitarian benchmark with no discrimination.

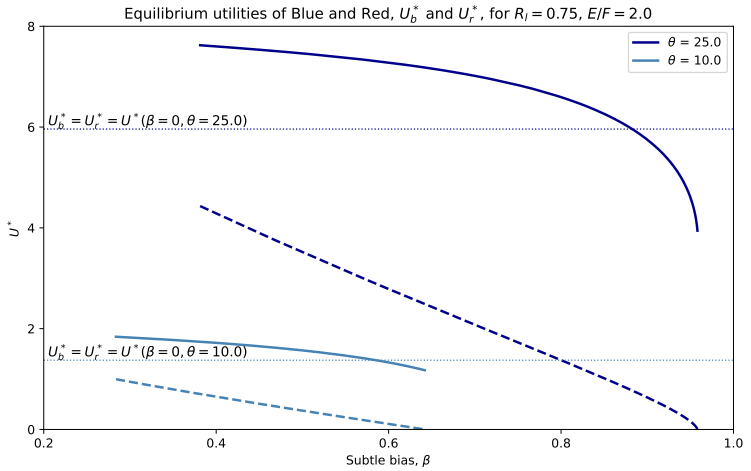


Figure 10: Utilities of Blue and Red workers, U_b^* and U_r^*.

Note: This figure presents the equilibrium utilities of Blue workers (solid lines) and Red workers (dashed lines), U_b^* and U_r^*, as functions of subtle bias, β, for $R_l = 0.75$, $E/F = 2.0$. The horizontal dotted lines represent the equilibrium utilities, $U_b^* = U_r^* = U^*$, in the benchmark case, $\beta = 0$, for $\theta = 25.0$ and $\theta = 10.0$.

Figure 11 shows profits as a function of β. In equilibrium, *profits increase with the bias*. Despite being inefficient, the biased equilibrium rigs the game in favor of firms by reducing workers' bargaining power. Thus, despite the deadweight losses, firms may prefer the biased equilibrium to the efficient no-discrimination equilibrium. Figure 11 shows that when β is sufficiently high, equilibrium profits are higher than in the efficient no-discrimination equilibrium (shown as horizontal lines). Intuitively, the subtle bias works as an implicit collusion device in the labor market by depressing red workers' outside options, allowing firms to capture a larger share of a shrinking surplus. Perhaps unexpectedly, the model can rationalize postmodern or Neo-Marxist theories that view discrimination against minorities as an implicit capitalist plot to increase profits by extracting more "surplus value" from workers. The surprising part is that we obtain this result in a model of perfect competition under rational expectations.

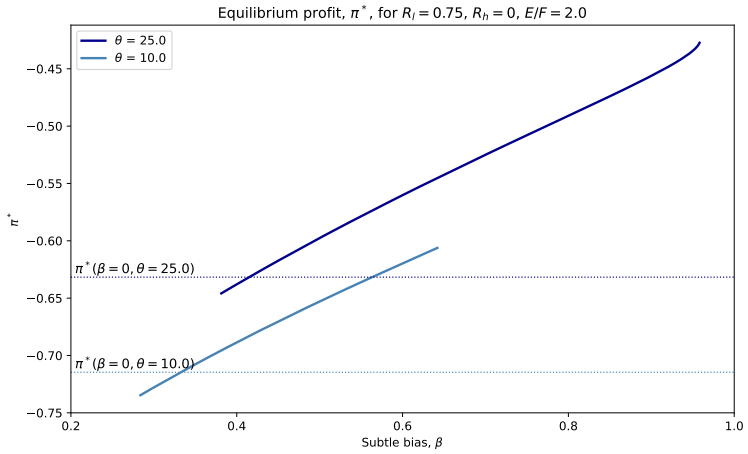


Figure 11: Firm profit, π^*.

Note: This figure presents the equilibrium profit of mixed firms, π^*, as a function of subtle bias, β, for $R_l = 0.75$, $R_h = 0$, $E/F = 2.0$. The horizontal dotted lines represent the equilibrium profit, π^*, in the benchmark case, $\beta = 0$, for $\theta = 25.0$ and $\theta = 10.0$.

4 Conclusion

This paper presents a model of promotions where vacancies, wages, and internal career paths are endogenously determined in a labor market equilibrium. Our findings highlight the impact of biased promotion practices on both workers and firms. Specifically, we show that a bias favoring blue workers leads to the emergence of two distinct types of firms: mixed firms and red firms. Mixed firms hire both blue and red workers, offer risky career paths, and pay higher wages, while red firms hire only red workers, offer stable, seniority-based promotions, and pay lower wages.

The model shows that a large bias against red workers is socially inefficient, reducing overall welfare for both blue and red workers. However, firms collectively benefit from this bias, as equilibrium profits increase with the level of bias. This result underscores the complex relationship between discrimination and economic efficiency, challenging the notion that meritocratic promotion systems are always beneficial for firms.

Our model generates several testable predictions that can guide future empirical research on promotions, discrimination, and inequality. These include the differential wage structures and promotion probabilities between mixed and red firms, the impact of bias on firm size and composition, and the overall effect on labor market dynamics. By providing a theoretical framework for understanding the consequences of biased promotions, this paper contributes to the broader literature on employer discrimination and firm performance.

References

AZMAT, G., V. CUÑAT, AND E. HENRY (2024): "Gender Promotion Gaps and Career Aspirations," *Management Science*.

BENSON, A., D. LI, AND K. SHUE (2024): ""Potential" and the gender promotion gap„" Working paper, University of Minnesota.

BERGSTROM, T. (1986): "Soldiers of fortune," *Equilibrium analysis: Essays in honor of Kenneth J. Arrow*, 2, 57–80.

BIRCAN, Ç., G. FRIEBEL, AND T. STAHL (2023): "Knowledge teams, careers, and gender„" Working paper, Goethe-University Frankfurt.

BOSQUET, C., P.-P. COMBES, AND C. GARCÍA-PEÑALOSA (2019): "Gender and promotions: Evidence from academic economists in France," *The Scandinavian Journal of Economics*, 121, 1020–1053.

CECCARELLI, M., C. HERPFER, AND S. ONGENA (2024): "Gender, performance, and promotion in the labor market for commercial bankers„" Swiss Finance Institute Research Paper.

CZIRAKI, P. AND A. ROBERTSON (2021): "Credentials matter, but only for men: Evidence from the S&P 500„" Working paper, University of Toronto.

FERREIRA, D. AND R. NIKOLOWA (2024): "Prestige, promotion, and pay," *The Journal of Finance*, 79, 505–540.

FREEMAN, S. (1996): "Equilibrium income inequality among identical agents," *Journal of Political Economy*, 104, 1047–1064.

HOFFMAN, M., L. B. KAHN, AND D. LI (2018): "Discretion in hiring," *The Quarterly Journal of Economics*, 133, 765–800.

HOSPIDO, L., L. LAEVEN, AND A. LAMO (2019): "The gender promotion gap: Evidence from central banking," *Review of Economics and Statistics*, 1–45.

HUANG, R., E. J. MAYER, AND D. P. MILLER (2023): "Gender bias in promotions: Evidence from financial institutions," *The Review of Financial Studies*, hhad079.

LAZEAR, E. P. AND S. ROSEN (1981): "Rank-order tournaments as optimum labor contracts," *Journal of Political Economy*, 89, 841–864.

PIKULINA, E. AND D. FERREIRA (2024): "Subtle discrimination," *Forthcoming in The Journal of Finance*.

ROSEN, S. (1986): "The theory of equalizing differences," *Handbook of Labor Economics*, 1, 641–6922.

——— (1997): "Manufactured inequality," *Journal of Labor Economics*, 15, 189–196.

WALDMAN, M. (1984): "Job assignments, signaling, and efficiency," *Rand Journal of Economics*, 15, 255–267.

——— (2013): "Classic promotion tournaments versus market-based tournaments," *International Journal of Industrial Organization*, 31, 198–210.

A Appendix

Proof of Proposition 1. By the definition of Θ, if $(R_l, \theta, E, F) \in \Theta$ at least one tight-labor-market equilibrium exists. Because $E > F$, in any equilibrium, there must be a positive mass of firms offering multiple vacancies, i.e., $n_\tau > 1$ for some positive mass of firms in $[0, F]$.

Consider one such equilibrium and let U^* denote the minimum lifetime utility implied by an equilibrium contract. Profit maximization implies that firms solve the problem:

$$\max_{n,w_l,w_h,\lambda,\mu} R_h - w_h + n(R_l - w_l) + \lambda\left(u(w_l) + \frac{1}{n}\theta u(w_h) - U^*\right) + \mu(n-1). \tag{A.1}$$

There must be firms with $n^* > 1$; we consider one such case and set $\mu = 0$. The problem is globally concave (see Ferreira and Nikolowa (2024)), and thus the first-order conditions determine a unique maximum given U^*. Thus, all firms must offer the same contract. The first-order conditions are

$$\frac{\partial L}{\partial w_l} = -n^* + \lambda^* u'(w_l^*) = 0 \tag{A.2}$$

$$\frac{\partial L}{\partial w_h} = -1 + \lambda^* \frac{1}{n^*}\theta u'(w_h^*) = 0 \tag{A.3}$$

$$\frac{\partial L}{\partial n} = R_l - w_l^* - \lambda^* \frac{1}{(n^*)^2}\theta u(w_h^*) = 0 \tag{A.4}$$

$$\frac{\partial L}{\partial \lambda} = u(w_l^*) + \frac{1}{n^*}\theta u(w_h^*) - U^* = 0 \tag{A.5}$$

Conditions (A.2) and (A.3) imply that the marginal utilities must be equal:

$$u'(w_l^*) = \theta u'(w_h^*). \tag{A.6}$$

Use (A.4) to isolate λ^* and replace it in (A.2) to find

$$R_l - w_l^* = \frac{\theta u(w_h^*)}{n^* u'(w_l^*)}. \tag{A.7}$$

Using (A.5) we can rewrite (A.7) as

$$R_l - w_l^* = \frac{U^* - u(w_l^*)}{u'(w_l^*)}. \tag{A.8}$$

Because the problem is globally concave, there is a unique solution for (A.8), $w_l(U^*)$. Because $u'(0) = \infty$, the left-hand side of (A.8) is greater than the right-hand side at $w_l^* = 0$, which implies $w_l'(U^*) < 0$. After finding w_l^*, we find w_h^* from (A.6), and then n^* from (A.5):

$$n^* = \frac{\theta u(w_h^*)}{U^* - u(w_l^*)} = \frac{E}{F}. \tag{A.9}$$

We find the equilibrium by solving for U^* in

$$u(w_l(U^*)) + \frac{F}{E}\theta u(w_h(U^*)) = U^*. \tag{A.10}$$

Because $w_l'(U^*) < 0$ and $w_h'(U^*) < 0$, there is only one solution to (A.10). □

Proof of Proposition 2. First, note that we must have $p_r(n, 0) = \frac{1}{n}$. Because Axiom 1 implies that $p_r(n, q)$ must be decreasing in q, without loss of generality, we can write $p_r(n, q) = \frac{1-\omega_r(n,q)}{n}$, where $\omega_r(n, q) \in [0, 1]$ is increasing in q and $\omega_r(n, 0) = 0$. Similarly, because $p_b(n, 1) = \frac{1}{n}$, we can write $p_b(n, q) = \frac{1+\omega_b(n,q)}{n}$, where $\omega_b(n, q) \in [0, n-1]$ is decreasing in q and $\omega_b(n, 1) = 0$. Axiom 2 then implies that $\omega_i(n, q) = \omega_i(q)$ where $i \in \{b, r\}$. Axiom 3 implies that $\omega_i(q)$ is linear. We then write $\omega_r(q) = q\omega_r$ (there is no intercept because $\omega_r(0) = 0$) and $\omega_b(q) = (1-q)\omega_b$ (because $\omega_b(1) = 0$). Probabilities must add up to 1: $q(1 + (1-q)\omega_b) + (1-q)(1-q\omega_r) = 1$, which implies $\omega_b = \omega_r := \beta$. Finally, $p_r(2, \frac{1}{2}) = \frac{1-\beta_b}{2} = \frac{1-\frac{1}{2}\beta}{2}$, so $\beta = 2\beta_b$. □

Proof of Proposition 3. The argument in the text implies that no firm can have a mixed workforce in equilibrium. Thus, firms must be fully segregated. "Blue firms" and "red firms" must have the same equilibrium profit. Because of full segregation, all workers in a firm face the same probability of promotion, which is $1/n_\tau^*$. Thus, they must be offered the unique contract described in Proposition 1. □

Proof of Lemma 1. Suppose C^* is an equilibrium with full segregation. If $U_r^*(C^*) \neq U_b^*(C^*)$, firms that employ workers with the highest utility have lower profits, which is not optimal

for them since they could target the low-utility workers instead. If $U_r^*(C^*) = U_b^*(C^*)$, all firms face the same maximization problem. Thus, the optimal number of workers must be unique, n^*. If $n^* = 1$, they all choose a safe career path, but this cannot be an equilibrium because $E > F$. Therefore, all firms must offer the same contract and $n^* > 1$. In this case, if firms are segregated, blue workers would prefer firms where $q(c_\tau^*) = 0$. Thus, the matching mechanism should assign some blue workers to these firms, which is a contradiction. □

Proof of Lemma 2. Part 1 follows from Lemma 1. To prove Part 2, note that if $c'' \neq c'$, we must have $q(c'') \neq q(c')$ because, otherwise, $\pi(c') \neq \pi(c'')$, which implies that one of these contracts is not optimal. But, if $q(c'') \neq q(c')$ and $n'' > 1$, profit is decreasing in q, which implies that $\pi(c') \neq \pi(c'')$. □

Proof of Lemma 3. Suppose there is a contract c'' such that $n'' = 1$. Then, $U_b(c'') = U_r(c'')$. Lemma 2 implies that there must exist another equilibrium contract c', which both types of workers accept, such that $n' > 1$. If blue workers also accept c'', then $U_b(c') = U_b(c'')$, and $U_r(c') \geq U_r(c'') = U_b(c'')$, implying that $U_r(c') \geq U_b(c')$, which is not possible; in a mixed firm, red workers must have a strictly lower utility than blue workers because $p_b(n', q(n')) > p_r(n', q(n'))$ for $q(n') \in (0, 1)$. □

Proof of Lemma 4. One of (14) and (15) must bind, otherwise, a firm can always increase n, increasing profits. Suppose the blue constraint (15) binds. Consider first a deviation contract c^d that, given $q^* > 0$, maximizes profit while ignoring the blue constraint but respecting the red constraint. Thus, by definition, contract c^d must violate the blue constraint, implying $U_b(c^d; q^*) < U_b^*$, where $U_b(c; q)$ denotes the utility of a blue agent that accepts contract c under belief q. Because this is an "optimal deviation," it must bind the red constraint, implying $U_r(c^d; q^*) = U_r^*$. Under c^d and q^*, the matching mechanism would thus assign only red workers to the deviating firm, which then implies that $q(c^d) = q^* > 0$ is not a rational belief. If only red workers are matched to the deviating firm, the only rational belief is $q(c^d) = 0$.

Thus, consider now a different optimal deviation, $c^{d'}$, assuming $q(c^{d'}) = 0$. Because the deviation is optimal, the red constraint must again bind, implying $U_r(c^{d'}; 0) = U_r^*$. We then have

$$U_b(c^{d'}; 0) = u(w_l^{d'}) + \frac{1+\beta}{n^{d'}}\theta u(w_h^{d'}) = U_r^* + \beta\theta\frac{u(w_h^{d'})}{n^{d'}}, \tag{A.11}$$

where the last step follows from $U_r(c^{d'}; 0) = U_r^*$. In general, for a belief function $q(.)$, we can write $U_b(c^d; q)$, where $q = q(c^d)$, for an optimal deviation c^d as

$$U_b(c^d; q) = U_r^* + \beta\theta\frac{u(w_h^d)}{n^d} = U_r^* + \beta\frac{R_l - w_l^d}{(1-q\beta)}u'(w_l^d), \tag{A.12}$$

where the last step follows from the first-order conditions. When the red constraint binds, w_l^d is independent of q. Thus, $U_b(c^d; q)$ increases with q, which implies $U_b(c^{d'}; 0) < U_b(c^d; q^*) < U_b^*$. Thus, $q(c^{d'}) = 0$ is the only rational belief under deviation $c^{d'}$, implying that contract c^* (for which the blue constraint binds) cannot be an equilibrium. □

Proof of Lemma 5. Lemma 4 implies that any deviation contract with $c^d \notin \{c_m^*, c_r^*\}$ that meets the red constraints is sub-optimal. Thus, even if there are rational beliefs that sustain such deviations, they are not profitable.

Suppose now that c^d violates both participation constraints at belief q^*. In this case, the matching mechanism will assign no one to this contract, implying that any q is rational for this deviation. Thus, we set $q(c) = q^*$ for any contract that attracts no worker. It follows that the profit is lower in this deviation than in the equilibrium contract.

The only remaining case is a deviation that, under q^*, violates the red constraint but not the blue constraint. In this case, the firm offers a risky career path ($n^d > 1$) only to blue workers. Under this deviation, the only rational belief is $q = 1$. Note that if a contract violates the red constraint at q^*, it will also violate it at $q = 1$. Let c^d denote the optimal contract subject to $U_b(c; 1) \geq U_b^*$. If $\pi(c^d) \leq \pi_m^*$, then the firm does not want to deviate to this contract. □