Biased Promotions*

Daniel Ferreira Radoslawa Nikolowa Elena S. Pikulina

February 2026

Abstract

We develop a competitive equilibrium model of biased promotions in which heterogeneity in firm size, wages, and career risk emerges endogenously from ex ante identical firms. Workers value wages and job amenities and differ only by a nonproductive label, "Blue" or "Red," while firms favor Blue workers in promotion decisions. The equilibrium features partial segregation through worker self-sorting: large, high-wage firms attract both groups and offer risky, biased promotion paths, whereas small, low-wage firms attract only Red workers and provide stable careers. Although firms individually prefer unbiased promotions, bias weakens workers' outside options, increasing industry profits while reducing efficiency. The model provides a competitive-market explanation for persistent group differences in promotions and earnings.

Keywords: Promotions, Discrimination, Inequality, Wages, Careers

*We thank Ashwini Agrawal, Eric French, Lubomir Litov, Michael Powell, David Webb, seminar participants at LSE, QMUL, São Paulo School of Economics (FGV), Cambridge University, HEC Montréal, University of Mannheim, Goethe University Frankfurt, Frankfurt School of Finance and Management, University of Calgary, and participants at the Asian Conference on Organizational Economics at HKU, the BSE Summer Forum Workshop on Tournaments, the Venice Workshop on Labor and Finance, the RAPS/RCFS Europe Meeting, the 2025 FMA Annual Meeting, the NBER Organizational Economics Meeting, the Queensland Corporate Finance Conference, the CEPR Paris Symposium, and the conference in honour of Ghazala Azmat. An earlier version circulated under the title "Promotion-Driven Inequality." 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 among central bankers (Hospido, Laeven, and Lamo (2019)), academic economists (Bosquet et al. (2019)), board directors (Cziraki and Robertson (2021)), lawyers (Azmat, Cuñat, and Henry (2025)), and bankers (Bircan, Friebel, and Stahl (2023); Huang, Mayer, and Miller (2023); Ceccarelli, Herpfer, and Ongena (2024)), among others. Much of the 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, yet men are promoted more often based on subjective assessments of greater "potential." Minni, Nguyen, and Sarsons (2025) find that managers from countries with more progressive gender attitudes promote women at higher rates. Cullen and Perez-Truglia (2023) show that social interactions with managers can explain about a third of the gender gap in promotions. Although promotion gaps are pervasive and economically consequential, accounting for a substantial share of the gender earnings gap,[1] their broader labor market consequences remain largely unexplored.

In this paper, we examine how biased promotion decisions shape labor market equilibrium outcomes. Specifically, we ask: How do these biases affect firm-level wage structures? Do they affect labor demand, firm size, and economic efficiency? Do they distort labor supply decisions, in particular the career paths workers pursue? Do they contribute to labor market segmentation? Answering these questions requires a theoretical framework that captures how promotion biases influence both firm behavior and workers' career paths.

To address these questions, we present a model of promotions where workforce size, wages, and internal career paths are determined in a competitive labor market equilibrium. A key feature of our model is that top positions carry greater prestige. As Baker, Jensen, and Murphy (1988) note, employees "*value the pay and prestige associated with a higher rank in the organization*" (p. 599, emphasis added). Our model takes this idea seriously: we explicitly model employees' preferences over pay and prestige.

As our benchmark, following Ferreira and Nikolowa (2024), we develop a compensating differentials model à la Rosen (1986), in which some jobs are perceived as being of higher quality than others. We can think of "high-quality jobs" as being more prestigious, or having other valuable amenities. A firm is endowed with a limited number of high-

[1] Cullen and Perez-Truglia (2023) speak of "*a consensus that this gap is primarily due to differences in promotion rates*" (p. 1708). Recent evidence in Arellano-Bover et al. (2024) is also consistent with this view.

quality jobs. Because jobs are indivisible and there are more workers than high-quality jobs, workers must be assigned to high-quality jobs through some randomization mechanism. We assume that firms cannot (literally) sell lottery tickets to job applicants. We then show that a firm may use the following alternative randomization mechanism: It creates several entry-level positions in a "low-quality job," fills them with young workers, and then randomly promotes a subset of them to the high-quality jobs. In practice, this randomness may operate through uneven opportunities to build promotable skills: access to promotable tasks, high-visibility projects, or mentoring. It may also stem from within-firm contests that are meritocratic yet uncertain in outcomes. What matters is that *ex ante* workers view promotions as random. We call this mechanism a *promotion lottery*.

A *promotion-lottery contract* is a long-term employment contract with limited promotion opportunities. It implies a well-defined career path featuring several properties of real-world internal labor markets, such as (among others) (i) the existence of "ports of entry," (ii) a preference for insiders in promotions, (iii) wages attached to jobs, (iv) rationing of top-level jobs, and (v) a positive correlation between rank and compensation.[2] In the model, firms compete for workers by offering promotion-lottery contracts in a perfectly competitive manner. Although promotion-lottery contracts increase the risk borne by risk-averse workers, we show that they represent a Pareto improvement over deterministic employment contracts.[3]

After presenting our benchmark model, we modify it to introduce biased promotions. Following Pikulina and Ferreira (2024), we assume that workers differ only by a payoff-irrelevant label, "Blue" or "Red," yet firms favor Blue workers when choosing among equally qualified promotion candidates. This form of discrimination is "subtle": it cannot be objectively verified and has no direct payoff consequences to firms. We interpret subtle discrimination as a form of friction: firms are unable to commit to promoting Blue and Red workers with equal probability. In practice, subtle discrimination may take the form of Red workers receiving fewer promotable assignments (Babcock et al. (2017); Bircan et al. (2024)) or having less access to mentoring (Athey et al. (2000); Matsa and Miller (2011); Kunze and Miller (2017)). We treat such biases as exogenously given and abstract

[2]Most of the theoretical literature explains the existence of internal labor markets as a mechanism for incentivizing workers (as in Lazear and Rosen (1981)) or matching skills to tasks (as in Waldman (1984)). In contrast, our model shows that internal labor markets may be optimal even in the absence of significant informational problems or worker heterogeneity.

[3]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)).

from belief updating (Bohren et al. (2019)).

Partly due to the threat of legal action, "overt" discrimination has become relatively rare, making subtle discrimination a more relevant phenomenon. Bennedsen et al. (2022) show that transparency about gender gaps increases women's promotion probabilities, suggesting a role for subtle biases in explaining gender promotion gaps. Consistent with our modeling of subtle discrimination, Huber, Lindenthal, and Waldinger (2021) and Ronchi and Smith (2021) show that firms may engage in discrimination without immediate profit consequences when workers are closely substitutable in terms of their skills.

We first consider a version of the model in which firms are subtly biased in promotion decisions, yet can offer different employment contracts to Blue and Red workers. These overtly discriminatory contracts may offset some of the distortions created by biased promotion decisions. We show that, in equilibrium, firms become fully segregated, hiring only Blue or only Red agents. While this equilibrium is efficient and egalitarian, both segregation and discriminatory contracts are likely to be socially unacceptable for reasons—such as externalities and moral constraints—that we intentionally leave outside the model.

In our main model, we rule out overtly discriminatory contracts. We then show that when firms are sufficiently biased toward promoting Blue workers, the equilibrium displays *partial segregation* of worker types. Two types of firms endogenously emerge. *Mixed firms* hire both Blue and Red workers. These firms create internal labor markets with *risky career paths*: they hire several entry-level workers, but promote only some of them to the top-level positions. Promoted workers enjoy a large promotion premium. Blue workers constitute the majority in Mixed firms and have a higher probability of promotion than Red workers (i.e., there is a Blue-Red promotion gap). In contrast, *Red firms* hire only Red workers. They offer *safe career paths*, where employees are promoted based on seniority. Red firms are smaller and pay lower entry-level wages than Mixed firms.

To identify additional equilibrium features, we numerically solve the model. We find that Red firms offer lower wages than Mixed firms at all hierarchical levels. Thus, in equilibrium, high- and low-wage firms coexist. Moreover, the unfavored group (Red) is less likely to work in high-wage firms, and even when they do, their expected wages are lower due to biased promotion decisions. Our model therefore offers a new explanation for the evidence that women are underrepresented in high-wage firms, and earn less than men in such firms (Bertrand and Hallock (2001); Card, Cardoso, and Kline (2016); Card et al. (2018); Goldin et al. (2017); Huneeus et al. (2021)). Crucially, we can rationalize these findings in a model where firm heterogeneity—the existence of high- and low-wage

firms—arises endogenously.

We also find that wages in Red firms decrease with the promotion bias, while wages in Mixed firms *increase* with the bias. Consequently, the wage gap between the two types of firms widens as the bias grows. Moreover, our analysis reveals that the favored group (Blue) may paradoxically be harmed by a bias in its favor. A larger bias weakens the bargaining position of workers, thus giving firms the upper hand in the labor market. Consistent with this result, we find that *profits increase with the bias*. Intuitively, the bias functions as an implicit collusion mechanism in the labor market, depressing the outside options of Red workers and allowing firms to capture a larger share of the surplus. Unexpectedly, our model offers a potential rationalization of some postmodern perspectives that argue that discrimination in the labor market causes worker fragmentation, potentially limiting collective bargaining power. Notably, this outcome emerges under perfect competition and rational expectations.

Our model predicts that when all firms are biased, *industry* profits increase with bias intensity, but overall efficiency declines. At the same time, an individual firm may find it optimal to deviate from industry norms and make costly investments to reduce its own bias. The Internet Appendix shows that when firms can invest in bias reduction, multiple equilibria may arise, including equilibria in which no firm invests. This result arises because a firm's benefit from eliminating its bias depends on how many other firms do the same, creating strategic complementarities.

Our model of promotions is based on Ferreira and Nikolowa (2024), who show the optimality of promotion-lottery contracts under fairly general conditions.[4] Following Pikulina and Ferreira (2024), we model discrimination as a consequence of subtle biases, which are defined as biases that operate only in "tie-breaking" decisions, such as the decision to promote one employee among several equally qualified candidates. Modeling subtle discrimination in our context is theoretically challenging because it requires specifying firms' and workers' beliefs about the shares of each worker type that apply to each type of firm. To pin down beliefs, our equilibrium concept imposes rational expectations on the equilibrium path, and rationalizable and "perfectly competitive" beliefs off the equilibrium path.

Our paper is related to the literature on promotions and internal labor markets. The concept of internal labor markets is due to Doeringer and Piore (1971). Baker, Gibbs, and

[4] Auriol and Renault (2008) and Auriol, Friebel, and Von Bieberstein (2016) develop related asymmetric-information promotion models in which promoted workers enjoy an increase in status. Although they do not discuss job creation, promotion lotteries can also be optimal in their models.

Holmstrom (1994) provide some evidence of internal labor markets from the records of a single firm. More recently, Huitfeldt et al. (2023) provide large-scale evidence of internal labor markets across several industries in Norway. They find evidence that firms have multiple internal labor markets with well-defined ports of entry, insider bias in promotions, and a strong correlation between rank and wages. Using the same methodology, Ewens and Giroud (2024) find evidence of internal labor markets in large US firms, characterized by pyramidal hierarchies and slot constraints on promotions. See also Lazear, Shaw, and Stanton (2018) for a theory and evidence on the importance of competition for job slots.

The theoretical literature on promotions and internal labor markets is extensive. Classic works include Lazear and Rosen (1981), Waldman (1984), Prendergast (1993), and Gibbons and Waldman (1999), while more recent contributions include dynamic moral hazard models by Axelson and Bond (2015) and Ke, Li, and Powell (2018). Most of this work emphasizes moral hazard, selection, or matching as explanations for promotion practices.

Incentives, selection, and matching may not suffice to explain all observed patterns associated with promotions. For example, Baker, Jensen, and Murphy (1988) highlight the widespread use of "job evaluation systems" that tie wage levels to "*(...) the "value of a job" according to a set of criteria such as the amount of training and education required, the total budget involved, the number of people supervised, and the amount of "independent decision-making" the job entails*" (p. 597). In other words, at least in some cases, wages appear to reflect *job attributes*, such as importance, responsibility, and difficulty, rather than worker characteristics, such as preferences and outside options, as traditional theories suggest.

Fahn and Klein (2025) also emphasize the need for alternative theories of promotions, especially in light of recent empirical puzzles, such as the evidence on the "Peter Principle" (Benson, Li, and Shue (2019)). Similar to our work, they present a model in which firms design promotion practices to optimally extract workers' surpluses. Their model is based on the exploitation of workers' overconfidence. In contrast, in this paper (as in Ferreira and Nikolowa (2024, 2025)), firms design contracts to optimally exploit workers' combined preferences for wages and job amenities.

Our theory of internal labor markets differs from previous works in two key respects. First, it demonstrates that internal labor markets may be optimal even in the absence of strong informational frictions or worker heterogeneity, which helps explain their prevalence across industries. Second, our model requires only minimal frictions, enabling us to study biased promotions in a (perfectly) competitive setting.

In addition to the evidence of biased promotions reviewed above, there is abundant evidence of managerial biases in other personnel decisions. Hoffman, Kahn, and Li (2018) show that biases, rather than superior information, account for the majority of instances in which managers exercise discretion to select candidates. Using a comprehensive audit study, Kline, Rose, and Walters (2022) document evidence of subtle forms of racial discrimination in hiring by large firms. While audit studies can precisely estimate average firm-specific biases, unique instances of subtle discrimination remain essentially undetectable. However, it is unclear whether audit studies can be used to detect biases in internal promotions.

This paper is also related to the extensive literature on discrimination in labor markets (see Fang and Moro (2011), Lang and Lehmann (2012), and Onuchic (2023) for reviews). Closely related is the seminal work of Lazear and Rosen (1990), who provide the first theoretical analysis of promotion gaps, in a model where men and women have exogenously different outside options. In contrast, in our model, Blue and Red workers have *endogenously* different outside options in equilibrium *because* of discrimination. Finally, our focus on the aggregate consequences of small biases is related to theoretical works on bias amplification, such as Lang, Manove, and Dickens (2005), Bartoš et al. (2016), Davies, Van Wesep, and Waters (2024), Siniscalchi and Veronesi (2021), and Onuchic and Ray (2023), among others.

In Section 2, we present a simplified version of the model. Section 3 presents the full model. Section 4 concludes. All proofs are provided in the Appendix. The Internet Appendix provides extensive analyses of omitted cases and extensions, particularly those related to firm entry and investments in bias reduction.

2 A Model of Biased Promotions

In this section, we present a simplified version of the model. To keep the analysis simple, we assume there is only one firm, with a fixed size. As a consequence of this assumption, the model in this section cannot generate our main results. Instead, the goal of this section is to explain the economic forces and underlying logic of the model. We present the full version of the model in Section 3, where we allow for a large number of firms to compete for workers, with no exogenous constraints on the number of workers they can hire.

For readability, in this section, we keep the presentation informal and discuss all proofs in the text. We postpone a more rigorous and formal analysis to Section 3.

2.1 A Simple Theory of Promotions

A firm may have several jobs $j \in J$, which are tasks that need to be performed by workers. Each job j has a quality attribute θ_j, capturing status, responsibility, autonomy, working conditions, or other positive nonpecuniary amenities. Workers care about job quality and wages: $u(\theta_j, w_j)$. For simplicity, we set $u(\theta_j, w_j) = \theta_j u(w_j)$, with $\theta_j > 0$, $u(w_j) > 0$, $u'(w_j) > 0$ and $u''(w_j) < 0$ for all $w_j > 0$, and $u'(0) = \infty$.[5] Note that the marginal utility of income increases with the nonpecuniary attribute, which is a standard assumption in the literature on status (see, e.g., Becker, Murphy, and Werning (2005)). This property is crucial for our main results.

Consider an economy with two identical workers who live for two periods. There are two jobs in this economy. The first job, the *standard job*, is in excess supply, has attribute $\theta_s = 1$, and pays a fixed wage $\underline{w}$ per period. We may think of this "job" as self-employment. The second job, the *high-quality job*, is in short supply; there is one firm that can offer one slot of this job per period. This job has attribute $\theta_h > 1$.

Suppose the firm lives forever and, at each period, a new cohort of two young workers enters the labor market, replacing the two workers who retire. The firm can hire only one worker per period, in which case its per-period profit is $\pi = R_h - w_h$, where $R_h > 0$ and w_h is the wage. Suppose the firm sets w_h at the beginning of a period, and workers apply for the position. The firm chooses w_h to fill the vacancy at minimum cost:

$$\theta_h u(w_h^*) = u(\underline{w}). \tag{1}$$

If more than one worker applies, the firm chooses one at random. We call this the *spot-wage contract* equilibrium. We may think of the selected worker as someone who is "promoted" from the standard job to the high-quality job.

We contend that a theory of promotions must explain at least four basic facts: (i) wages increase upon promotion; (ii) workers strictly prefer to be promoted (i.e., there are "promotion rents"); (iii) slot constraints imply rationing of "high-quality" jobs; and (iv) firms create internal labor markets. This spot-contract theory fails to account for all four of the required facts, as wages decrease upon promotion ($w_h^* < \underline{w}$), promoted workers receive no rents, and there is no rationing or internal labor markets.

There is an additional problem with this theory: The equilibrium implied by (1) is

[5] 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)).

inefficient. To see this, notice that a worker's expected per-period utility is

$$\frac{1}{2}\theta_h u(w_h^*) + \frac{1}{2}u(\underline{w}) = u(\underline{w}). \tag{2}$$

Because $\theta_h u'(w_h^*) > u'(\underline{w})$, if a (hypothetical) social planner changes wages to $w_h^* + \Delta$ and $\underline{w} - \Delta$ for $\Delta > 0$ small, the aggregate wage bill is unchanged, and the workers' expected utility becomes

$$\frac{1}{2}\theta_h u(w_h^* + \Delta) + \frac{1}{2}u(\underline{w} - \Delta) > u(\underline{w}). \tag{3}$$

Thus, this is a Pareto improvement, implying that choosing w_h^* is socially inefficient. The source of inefficiency is that *marginal utilities are not equalized* when (1) holds. Note that this happens because $\theta_h > 1$ (i.e., the high-quality job is more desirable than the standard job).

The possibility of a Pareto improvement implies that the firm is not maximizing its profit. To consider an alternative solution, suppose the firm sells "lottery tickets" to both workers for a fee. The firm hires the winner of the lottery for a wage $\hat{w}_h$, and the loser works in the standard job. Without loss of generality, assume that only the loser pays a fee, f (equivalently, we can interpret $\hat{w}_h$ as the "winning" wage net of the fee.) The workers are willing to enter the lottery if

$$\frac{1}{2}\theta_h u(\hat{w}_h) + \frac{1}{2}u(\underline{w} - f) \geq u(\underline{w}). \tag{4}$$

The firm chooses $(f, \hat{w}_h)$ to maximize its profit, $\hat{\pi} = R_h - \hat{w}_h + f$, subject to the participation constraint in (4). The first-order conditions imply

$$u'(\underline{w} - f^*) = \theta_h u'(\hat{w}_h^*). \tag{5}$$

The solution $(\hat{w}_h^*, f^*)$ must satisfy (4) (with equality) and (5). We call this the *spot-lottery contract* equilibrium. Because the firm may choose $f = 0$ if it wishes, the firm must do at least as well under spot lotteries as it does under spot wages. Since setting $f = 0$ in (5) implies $\hat{w}_h > \underline{w}$, the participation constraint is slack at this wage, and the firm can increase its profit by charging a strictly positive fee. Thus, we must have $f^* > 0$.

At first blush, it might seem puzzling that risk-averse workers are willing to pay for lotteries with negative net expected values. The reason for this behavior is that the attribute $\theta_h > 1$ effectively makes workers behave *as if* they are risk lovers. Formally, utilities are state-dependent, and the marginal utility is higher for all wage levels in state h (i.e., when the worker is assigned to the high-quality job) than in self-employment. The spot-lottery contract exploits this convexity, thereby increasing the firm's profits. Crucially important

is the assumption that job assignments affect the marginal utility of income.[6]

The spot-lottery contract resembles a promotion tournament. Crucially, wages increase upon "promotion" ($\hat{w}_h^* > \underline{w}$), promoted workers enjoy rents ($\theta_h u(\hat{w}_h^*) > u(\underline{w})$), and high-quality jobs are rationed. However, we still lack an internal labor market.

The practical relevance of spot-lottery contracts is questionable. While we often see modest application fees for positions, it is unlikely that such fees can be sufficient to compensate firms for raising their wages to levels above those in the outside market. The primary practical challenge in implementing spot lotteries is that firms may struggle to collect substantial fees from liquidity-constrained applicants. In our simple model, this could arise because the outside wage $\underline{w}$ cannot be pledged, either because of legal constraints (e.g., the worker would need to pledge a significant share of her future income to pay fees for unsuccessful job applications) or because $\underline{w}$ is not a wage but the money-equivalent of a bundle of pecuniary and nonpecuniary benefits of self-employment.

If firms cannot charge substantial application fees, they may still use their productive technologies to extract more surplus from workers who value the high-quality job. Suppose that the firm may create (at no cost) a new job, which we call job l. We assume that this job has attribute $\theta_l = 1$ (as the standard job). The firm can create as many positions as it wants for this job. Each filled job-l position yields revenue R_l, but we assume $u(R_l) < u(\underline{w})$, so these jobs are individually inefficient. This implies that the firm does not create job-l positions under either spot-wage contracts or spot-lottery contracts. However, the firm now has the option of offering the following long-term employment contract to each young worker: "You work in job l for wage $\tilde{w}_l$ while young. When old, with probability 0.5 you will be promoted to job h and earn wage $\tilde{w}_h$." We call this contract a *promotion-lottery contract*. Essentially, this contract creates an internal labor market for the two workers, with an "up-or-out" structure; the worker who is not promoted will prefer to leave the firm and earn $\underline{w}$ in the standard job.

The workers will accept the promotion lottery contract if (assuming no discounting)

$$u(\tilde{w}_l) + \frac{1}{2}\theta_h u(\tilde{w}_h) + \frac{1}{2}u(\underline{w}) \geq 2u(\underline{w}). \tag{6}$$

The firm chooses $(\tilde{w}_l, \tilde{w}_h)$ to maximize its profit, $\tilde{\pi} = R_h - \tilde{w}_h + 2(R_l - \tilde{w}_l)$, subject to the

[6] Auriol and Renault (2008), Auriol, Friebel, and Von Bieberstein (2016), and Ferreira and Nikolowa (2024) make this assumption in models of promotions. The literature on status similarly assumes that status increases the marginal utility of income; see Hopkins and Kornienko (2004), Becker, Murphy, and Werning (2005), and Ray and Robson (2012), among others.

participation constraint in (6). The first-order conditions imply

$$u'(\tilde{w}_l^*) = \theta_h u'(\tilde{w}_h^*). \tag{7}$$

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. 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 high-skilled workers who have been promoted. By contrast, in promotion-lottery models, workers are (productively) identical, and there is no non-contractible effort. Promotion premiums exist because cost-minimizing contracts must equalize workers' marginal utilities across jobs. Intuitively, if marginal utilities differ across jobs, workers are willing to pay a price to transfer wages from the job with lower marginal utility to the job with higher marginal utility.

Contract $(\tilde{w}_l^*, \tilde{w}_h^*)$ is still inefficient relative to $(f^*, \hat{w}_h^*)$ because working in job l is an inefficient way to pay for the promotion lottery. However, if charging fees is infeasible, the promotion lottery is the second-best contract, as long as the profit under promotion lotteries is larger than the profit under spot wages: if $\tilde{\pi}^* \geq R_h - w_h^* \Leftrightarrow \tilde{w}_h^* - w_h^* \leq 2(R_l - \tilde{w}_l^*)$.[7]

Promotion-lottery contracts exhibit several key properties of internal labor markets observed in the real world, including a positive correlation between pay and rank, insider bias in promotions, rationing of top-level jobs, and the presence of ports of entry, all without invoking asymmetric information or learning. The lottery element of promotions can be implemented with either objective or subjective performance criteria. Promotions based on objective performance can be perceived as a lottery due to the stochastic nature of production and incomplete information about one's competitors (and sometimes of one's own type). Similarly, subjectively-assessed promotions are akin to a lottery if the assessor's information set and biases are unknown.

Promotion lotteries offer a "third view" of internal labor markets. When firms control access to high-quality jobs, they design internal career paths leading to these jobs and create low-quality jobs that function as ports of entry. These entry-level jobs serve as currency for workers to pay for a chance to compete for a promotion.

[7] In a more general model, Ferreira and Nikolowa (2024) show that this condition always holds for θ_h sufficiently high.

2.2 Biased Promotions

Suppose now that the two workers have different labels, $i \in \{b, r\}$, for Blue and Red. As in Pikulina and Ferreira (2024), labels are observable and productively irrelevant. The firm offers a promotion-lottery contract to each worker. 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. We assume that subtle biases are exogenously given and fixed, that is, we abstract from belief updating (Bohren et al. (2019)).

The existence of a subtle bias implies that the firm cannot commit to a fair promotion lottery. In addition, we assume that the firm cannot *overtly discriminate*: wages $(\tilde{w}_l, \tilde{w}_h)$ must be the same for both types of workers. Now, the lifetime participation constraints are type-dependent:

$$u(\tilde{w}_l) + \frac{1+\beta_b}{2}\theta_h u(\tilde{w}_h) + \frac{1-\beta_b}{2}u(\underline{w}) \geq 2u(\underline{w}) \tag{8}$$

for the Blue worker, and

$$u(\tilde{w}_l) + \frac{1-\beta_b}{2}\theta_h u(\tilde{w}_h) + \frac{1+\beta_b}{2}u(\underline{w}) \geq 2u(\underline{w}) \tag{9}$$

for the Red worker. To preserve the property that Red prefers the high-quality job, assume $(1-\beta_b)\theta_h > 1$. Because we must also have $\theta_h u(\tilde{w}_h) \geq u(\underline{w})$ in equilibrium (otherwise the firm is unable to retain any old worker in job h), we have that (9) implies (8). Thus, the firm chooses $(\tilde{w}_l, \tilde{w}_h)$ to maximize its profit subject to (9).

The first-order conditions imply

$$u'(\tilde{w}_l^*) = (1-\beta_b)\theta_h u'(\tilde{w}_h^*). \tag{10}$$

Note that, as before, wages increase upon promotion. However, unlike the unbiased 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 $1-\beta_b$. Intuitively, the bias β_b reduces the expected value of the wage in the top job for the Red worker. The firm chooses

its optimal contract to equalize the Red worker's marginal utilities *as perceived by her*.

Condition (10) implies that the equilibrium is *inefficient* (relative to the second-best, i.e., unbiased promotion lotteries). To see this, note that the firm's equilibrium profit is decreasing in $\beta_b \in [0,1]$,[8] which implies that the firm would prefer to be unbiased. Thus, we can think of β_b as a friction: firms with positive subtle biases are unable to commit to a fair promotion lottery. This friction implies a smaller *promotion premium* $(\tilde{w}_h^* - \tilde{w}_l^*)$ than the second-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," in the sense that the second-best requires a larger promotion premium.

We conclude that a subtle bias towards one type of worker can lead to flatter career wages (i.e., smaller promotion premiums), lower profits, and a less efficient employment contract. However, the simple model in this section has one main limitation: there is a single monopsonistic firm. In the next section, we extend the model to allow for multiple firms that compete for workers.

3 A Competitive Model of Biased Promotions

In this section, we extend the simple model from the previous section to consider the case of multiple firms competing in the same labor market. All proofs are in the Appendix. The footnotes provide further discussions on the assumptions.

3.1 Setup: Firms, Workers, and Contracts

Consider an economy with a mass F of identical, infinitely lived firms. Time is discrete. Each firm is indexed by $\tau \in [0,F]$. In each period, a firm may have vacancies in two jobs, h and l. Jobs are indivisible, and a worker can perform only one job per period. Firms have a fixed supply of slots in job h, which we set to 1. Firms can also employ a mass $n \geq 0$ of workers in job l. There are no slot constraints for job l; firms can create as many l-job slots as they wish (i.e., n is a choice variable). A firm receives revenue $R_j > 0$ per period for each unit mass of workers employed in job $j \in \{h,l\}$.[9]

[8] Differentiating the Lagrangian with respect to β_b yields $-\frac{\lambda^*}{2}[\theta_h u(\tilde{w}_h^*) - u(\underline{w})] < 0$, where λ is the (positive) multiplier.

[9] The technology exhibits constant returns to labor; revenue from job l increases linearly with n at a constant marginal rate R_l. We choose this specification for simplicity. Our results can also be derived in an alternative model where the marginal productivity of labor, $R_l(n)$, is decreasing in n.

In each period, a mass $E > F$ of young agents enters the labor force and retires after two periods. Each agent has an observable and productively irrelevant label, $i \in \{b, r\}$, for Blue and Red. The proportion of type-i agents in the population is α_i, with $\alpha_b + \alpha_r = 1$. We assume that firms compete in the labor market by offering long-term employment contracts to young workers, which are designed as promotion-lottery contracts. In these contracts, young workers perform job l (the entry-level job) for one period. In the second period of employment, some workers are promoted to job h; those who fail to gain promotion leave the firm.[10] Firms commit to promoting only insiders but retain discretion over promotion decisions.[11] Firms attract job applicants by advertising (and committing to) a number of vacancies for job l, (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.[12] The total number of vacancies is $n = n_b + n_r \geq 1$.[13]

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 $i \in \{b, r\}$. If $n = 1$, the promotion lottery is degenerate; all entry-level workers are promoted to job h (which has a fixed supply of 1) after one period. Thus, if $n = 1$, 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's *per contract* profit (assuming no discounting) is thus

$$\pi(c) = R_h - w_h(c) + nR_l - n_b w_{lb} - n_r w_{lr}, \tag{11}$$

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

$$w_h(c) = \begin{cases} \min\{n_b, 1\} w_{hb} + \max\{1 - n_b, 0\} w_{hr} & \text{if} \quad w_{hb} \leq w_{hr} \\ \min\{n_r, 1\} w_{hr} + \max\{1 - n_r, 0\} w_{hb} & \text{if} \quad w_{hb} \geq w_{hr} \end{cases}. \tag{12}$$

Note that if $w_{hb} = w_{hr}$, the firm is indifferent as to whom to promote, and thus, any promo-

[10] 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.

[11] Our results remain unchanged as long as firms are more likely to promote insiders than outsiders.

[12] To streamline the analysis, we assume that a dismissed worker receives no compensation from the firm. While severance compensation can sometimes be optimal (as shown in Ferreira and Nikolowa (2024)), its availability does not affect the optimality of promotion lotteries or the qualitative features of the equilibrium.

[13] Choosing $n < 1$ is not optimal because the firm would leave some slots in job h vacant.

tion rule is rational.

As in the previous section, agents derive utility from money and job attributes. Job l has attribute $\theta_l = 1$ and job h has attribute $\theta_h = \theta > 1$. All agents—regardless of their labels—have the same preferences. We denote the (per period) utility of an unemployed (or self-employed) 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 of a promotion lottery leading to job h.[14] Without loss of generality, from now on we normalize the unemployment utility to zero: $\underline{u} = 0$.

We define a worker's lifetime utility under contract c as

$$U_i(c) = u(w_{li}) + p_i(c)\theta u(w_{hi}). \tag{13}$$

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 equilibrium probability of being promoted given contract c. Note that we impose that beliefs depend only on the contract c that applies to the worker, and not on other contracts $c' \neq c$ that might be available but were not chosen by the worker. Additionally, we assume that all workers and firms share the same set of beliefs.

3.2 Equilibrium Conditions

We model the matching process as arising from directed search. First, all firms simultaneously post (and commit to) contracts. Each firm can post a single contract. Then, all young workers apply for jobs. Workers and firms are infinitesimal and, thus, ignore strategic considerations (i.e., search is competitive, in the sense of Wright et al. (2021)).

Suppose firm τ posts contract $c_\tau = (n_{b\tau}, w_{lb\tau}, w_{hb\tau}, n_{r\tau}, w_{lr\tau}, w_{hr\tau})$, while the set of contracts posted by all other firms is $C_{-\tau}$. We denote the set of all contracts by $C = (c_\tau, C_{-\tau})$, and we refer to the "self-employment" contract as $\underline{c}$. Let $\mu_{i\tau}(c_\tau, C_{-\tau})$ be the

[14] The "inefficient job" assumption, 1(ii), is needed only because of the constant returns technology. Under a decreasing returns technology, firms would employ workers even if their marginal productivity is lower than the self-employment utility, because workers are willing to work in the entry-level job, hoping to be promoted.

mass of workers of type i who apply to firm τ. We can think of $\mu_{i\tau}(c_\tau, C_{-\tau})$ as firm τ's *residual labor supply function*, that is, the supply of type-i workers to firm τ given the contracts of all other firms. We impose two conditions on the residual supply functions.

Condition 1 (Symmetry). *If $c_\tau = c_{\tau'}$, then $\mu_{i\tau}(c_\tau, C_{-\tau}) = \mu_{i\tau'}(c_{\tau'}, C_{-\tau'})$ for all $i \in \{b,r\}$ and $\tau, \tau' \in [0,F]$.*

Condition 1 simply says that workers perceive two firms offering the same contract as equivalent. In other words, the firm index τ does not affect workers' decisions to apply for vacancies.

Define

$$U_i^*(C) = \max_{c \in C \cup \{\underline{c}\}} U_i(c). \tag{14}$$

That is, $U_i^*(C)$ is the maximum utility a type-i worker obtains when free to choose any contract in C or the self-employment contract, $\underline{c}$. The second condition we impose is as follows.

Condition 2 (Perfect Competition). *The residual labor supply functions are perfectly elastic:*

$$\mu_{i\tau}(c_\tau, C_{-\tau}) = \begin{cases} n_{i\tau}, & \text{if } U_i(c_\tau) \geq U_i^*(C_{-\tau}) \\ 0, & \text{if } U_i(c_\tau) < U_i^*(C_{-\tau}) \end{cases} \tag{15}$$

for all $\tau \in [0,F]$ and $i \in \{b,r\}$.

Intuitively, Condition 2 implies that firms are "price takers" in the labor market:[15] they view their residual supply of workers as perfectly elastic. All else being constant, if a firm offers a contract that is weakly better than the competition, it will fill all of its vacancies. Conversely, if a firm offers a slightly inferior contract than the competition, it attracts no workers.[16]

Because only under a tight labor market there is effective competition among firms, from now on, we restrict attention to *tight-labor-market equilibria*, that is, equilibria in which all workers are employed.[17] We define an equilibrium as follows.

[15] Or, more precisely, firms take the set of competing contracts, $C_{-\tau}$, as given.

[16] A possible microfoundation for Condition 2 is Bertrand competition in the labor market. In our model, firms have an optimal "capacity" (i.e., an optimal number of workers given market conditions); therefore, Bertrand competition does not imply zero profits. This is reminiscent of Kreps and Scheinkman's (1983) classic model of price competition under capacity constraints.

[17] In a similar setup, Ferreira and Nikolowa (2024) show parameter conditions such that a tight-labor-market equilibrium exists. More generally, an equilibrium always exists for any set of parameters, but it may involve some unemployed workers or idle firms. In a slack labor market, some or all Red workers will be unemployed, and the equilibrium utility of all Red workers will be zero, $\underline{u} = 0$.

Definition 1. *An equilibrium is a set of contracts* $C^* = \{c_\tau^* : \tau \in [0,F]\}$ *and residual labor supply functions* $\mu_{i\tau}(c_\tau, C_{-\tau})$ *satisfying Conditions 1 and 2 such that:*

1. *If* $c_\tau^* \in C^*$, *then* $c_\tau^* \in \arg\max_c \pi(c)$ *subject to* $U_i(c) \geq U_i^*(C_{-\tau}^*)$, *for* $i \in \{b,r\}$;
2. $n_{i\tau}^* = \mu_{i\tau}(c_\tau^*, C_{-\tau}^*)$, *for all* $\tau \in [0,F]$ *and* $i \in \{b,r\}$;
3. $\int_0^F n_{i\tau}^* d\tau \leq \alpha_i E$ *for* $i \in \{b,r\}$.

Part 1 states that firms choose contracts to maximize profits, subject to workers maximizing utility, assuming all other firms' contracts are given. Part 2 implies that each firm's demand for type i equals its residual labor supply. Part 3 requires aggregate labor demand to be no greater than aggregate labor supply. If Part 3 holds with equality, we have a *tight labor market* equilibrium: all workers are employed by the firms in the sector. If, for some $i \in \{b,r\}$, we have that Part 3 holds with strict inequality, we have a *slack labor market* equilibrium. In the latter case, we may have some unemployed workers of one or both types.

3.3 Benchmark: Unbiased firms

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

$$U(c) = u(w_l) + \frac{1}{n}\theta u(w_h), \tag{16}$$

and a firm's per-contract profit is

$$\pi(c) = R_h - w_h + n(R_l - w_l). \tag{17}$$

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).

Under the assumption that the parameters are such that a tight-labor-market equilibrium exists, the following proposition shows that a tight-labor-market equilibrium is unique and has all firms offering the same contract.

Proposition 1 (Benchmark Equilibrium Contract). *There is a unique tight-labor-market equilibrium contract* $c^* = (n^*, w_l^*, w_h^*)$, *which is offered by all firms.*

Three intuitive conditions determine the equilibrium (see the proof of Proposition 1 for the omitted steps). First, the marginal utilities in each job must equalize:

$$u'(w_l^*) = \theta u'(w_h^*). \tag{18}$$

This condition is the same as in the simple model of the previous section and implies that wages must increase upon promotion. Second, we have

$$R_l = w_l^* + \frac{1}{n^*} \frac{\theta u(w_h^*)}{u'(w_l^*)}. \tag{19}$$

This condition equates the marginal productivity of job l, R_l, to the marginal cost of hiring 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 (19) does not hold, firms would like to either increase or reduce the number of job positions they offer. Condition (19) 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 under a tight labor market imply $n^* = E/F$. 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^*.[18]

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 aligns 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 empirical prediction of the model. 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.[19]

[18] The figure shows wages only for parameter configurations for which a tight-labor market exists and $\ln w_h^* > 0$.

[19] Consistent with this prediction, Bidwell et al. (2015) show that starting compensation varies little across investment banks, but pay rises more quickly with career progression in more prestigious banks.

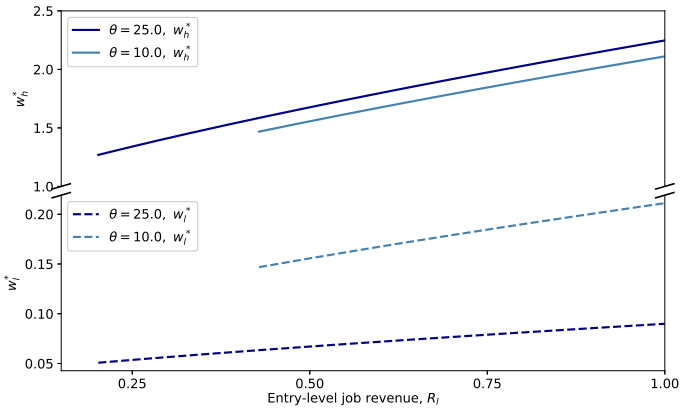


Figure 1: Equilibrium wages in the benchmark case.

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 $E/F = 2.0$

Overall, Figure 1 illustrates that promotion lotteries result in significant ex-post inequality among ex-ante identical workers. However, this is "good" inequality: promotion lotteries represent a Pareto improvement over deterministic employment contracts. Rosen (1997) calls this type of efficient lottery-driven inequality "manufactured inequality."

The equilibrium of this benchmark model is constrained-efficient. To see this, notice that the assignment of workers to job h must involve some kind of lottery, because the supply of h jobs is $F < E$, and jobs are indivisible. We know from the previous section that, in such a case, the efficient contract involves firms charging workers application fees. However, if firms cannot collect fees from job applicants, the second-best solution is to use promotion lotteries in which workers are initially assigned to job l. Conditional on the use of promotion lotteries, efficiency requires equalizing marginal utilities across the two jobs, as in (18), and equalizing marginal productivity to marginal cost, as in (19).

Suppose now that $\alpha_b < 1$. Since Blue and Red workers are productively identical, if the firm is unbiased, 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* towards one type, as in Subsection 2.2. We consider this possibility next.

3.4 Biased Promotion Probabilities

In this subsection, we extend Pikulina and Ferreira's (2024) notion of subtle discrimination to the case of $n = n_b + n_r$ agents. To make the analysis more intuitive, in this subsection, we interpret the number of workers, n, as an integer.

Consider a firm that is biased in favor of Blue workers. Let $q = \frac{n_b}{n}$ denote the proportion of Blue workers competing for the top job in the firm. Suppose a type-i worker faces probability $p_i(n,q)$ of being promoted, for $i \in \{b,r\}$. The probability that the firm promotes one of the n_i type-i workers is thus $n_i p_i(n,q)$. We define the firm's bias towards Blue workers as:

$$\beta_b(n,q) = n\left[qp_b(n,q) - (1-q)p_r(n,q)\right] - (2q-1). \tag{20}$$

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. In Subsection 2.2, we had $n = 2$, $q = 0.5$, $p_b = \frac{1+\beta_b}{2}$, and $p_r = \frac{1-\beta_b}{2}$, implying that the bias was $\beta_b(2,0.5) = \beta_b$ in that case.

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

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

Axiom 1 simply says that because the firm is biased towards Blue workers, 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 minus one 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 halves the probability of promotion for everyone. 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: $p_i(n,q_0+\varepsilon) - p_i(n,q_0)$ is independent of q_0 for any $\varepsilon < 1-q_0$.*

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 the function $p_i(n,q)$.

Proposition 2 (**Biased Promotion Probabilities**). *Axioms 1-3 imply that the individual promotion probabilities are*

$$p_b(n,q) = \frac{1+(1-q)\beta}{n} \text{ and } p_r(n,q) = \frac{1-q\beta}{n}, \tag{21}$$

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 by Axioms 1-3, they are consistent with multiple mechanisms for choosing whom to promote. For example, the following two-step procedure delivers (21): (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).[20]

3.5 Equilibrium with Overt Discrimination

In this subsection, we characterize the equilibrium when firms are allowed to "overtly discriminate," that is, firms may offer contracts that are exclusive to a type of worker, $i \in \{b,r\}$.

When firms are subtly biased, workers care about the identity of their coworkers, as they 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 are more likely to promote Blue workers. Thus, under the same contract, workers—both Blue and Red—prefer to work in firms with fewer Blue workers.

We first note that 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 a Red worker in a mixed-type firm has

[20] We also need some parameter restrictions for the probabilities in (21) to be between 0 and 1. We restrict β to be in $[0,1]$, implying $p_r(n,q) \in [0,1]$. For $p_b(n,q)$ to be lower than 1, we need $1-q \leq n-1$. In a discrete setting, we must have at least two people, implying $n \geq 2 \Rightarrow p_b(n,q) \leq 1$. However, for convenience, we later work with a continuous n, thus it is technically possible that $n \in (1,2)$. To streamline the exposition, we will henceforth ignore this case, but discuss it in the proofs where appropriate.

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 promotion.

Formally, let c_i^0 denote a contact offered to type i where $n_i = 0$. That is, a firm that offers c_i^0 overtly discriminates against type i by refusing to hire workers of that type. Let c_i^* denote the benchmark contract defined in Proposition 1 when offered to type i only. We then have the following result.

Proposition 3 (**Full Segregation under Overt Discrimination**). *If firms are subtly biased ($\beta > 0$) and overt discrimination is possible, there is a unique equilibrium where only two contracts are offered, $\left(c_b^*, c_r^0\right)$ and $\left(c_b^0, c_r^*\right)$. 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 consequence of this trend is that firms cannot use employment contracts to "de-bias" their promotion practices.

3.6 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 overtly discriminatory employment contracts. Our analysis here extends the model in Subsection 2.2 to multiple firms.

Each firm $\tau \in [0,F]$ must offer a non-discriminatory contract, $c_\tau = (n_\tau, w_{l\tau}, w_{h\tau})$. A firm's per-contract profit is $\pi(c_\tau) = R_h - w_{h\tau} + n_\tau(R_l - w_{l\tau})$. If more than n_τ workers apply to the firm, it must randomly choose among them, regardless of type. That is, we

assume that firms' hiring decisions are unbiased, unlike their promotion decisions.[21]

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 $q(c_\tau, C_{-\tau})$ denote the proportion of Blue workers firm τ expects to hire when it offers contract c_τ and the set of contracts offered by all other firms is $C_{-\tau}$. A type-i worker's lifetime utility under c_τ when all other firms offer $C_{-\tau}$ is now

$$U_i(c_\tau, C_{-\tau}) = u(w_{l\tau}) + p_i(n_\tau, q(c_\tau, C_{-\tau}))\theta u(w_{h\tau}), \tag{22}$$

where $p_i(n,q)$ is given by (21), and $c_\tau = (n_\tau, w_{l\tau}, w_{h\tau})$. Function $q(c_\tau, C_{-\tau})$ represents agents' *beliefs* about q_τ. We assume that all workers and firms share the same beliefs.

Let C^* denote an equilibrium set of contracts and $q(c^*_\tau, C^*_{-\tau})$, for $\tau \in [0,F]$, denote the beliefs on the equilibrium path. Let $q(c_\tau, C^*_{-\tau})$ denote the belief about q_τ given contracts $(c_\tau, C^*_{-\tau})$, where c_τ is any feasible (but not necessarily equilibrium) contract offered by firm $\tau \in [0,F]$. To pin down the full equilibrium set of beliefs, both on and off path, we need to impose some further structure. We first define *individually-rational beliefs* as follows.

Definition 2 (Individually-Rational Beliefs). *We say that belief $q(c_\tau, C^*_{-\tau})$ is **individually rational** for contracts $(c_\tau, C^*_{-\tau})$ if*

$$q(c_\tau, C^*_{-\tau}) = \begin{cases} 1 & \text{if} \quad U_b(c_\tau, C^*_{-\tau}) \geq U^*_b(C^*_{-\tau}) \text{ and } U_r(c_\tau, C^*_{-\tau}) < U^*_r(C^*_{-\tau}) \\ 0 & \quad\quad U_b(c_\tau, C^*_{-\tau}) < U^*_b(C^*_{-\tau}) \text{ and } U_r(c_\tau, C^*_{-\tau}) \geq U^*_r(C^*_{-\tau}) \end{cases} \tag{23}$$

In words, individual rationality requires beliefs about q_τ to be consistent with workers' *individual rationality* (i.e., participation) constraints. If Blue's IR constraint holds but Red's doesn't, agents expect only Blue workers to apply to contract c_τ; individual rationality thus requires $q = 1$. Similarly, if Red's IR constraint holds but Red's doesn't, agents expect only Red workers to apply to contract c_τ; individual rationality thus requires $q = 0$.

While it is reasonable to impose individual rationality on beliefs, it is a weak constraint: for any given set of contracts $(c_\tau, C^*_{-\tau})$, there is generally a large set of beliefs q_τ that are consistent with individual rationality. In particular, if both IR constraints hold or if both don't hold, individual rationality places no restrictions on q_τ. Thus, we impose two conditions on the set of beliefs. First, the next condition applies to beliefs on the equilibrium path.

[21] This assumption is made to streamline the presentation and is without loss of generality. Because we consider only tight-labor market equilibria, subtle biases in hiring have no impact on who is hired in equilibrium, as each firm receives the same number of applications as the number of vacancies it posts. Intuitively, subtle biases only matter when there is rationing of positions, as is the case with competitive promotions.

Condition 3 (Rational Expectations). *Let C^* denote an equilibrium set of contracts. Then, $\mu_{b\tau}(c^*_\tau, C^*_{-\tau}) = q(c^*_\tau, C^*_{-\tau})[\mu_{b\tau}(c^*_\tau, C^*_{-\tau}) + \mu_{r\tau}(c^*_\tau, C^*_{-\tau})]$ for all $\tau \in [0, F]$.*

Condition 3 requires agents to have rational expectations in equilibrium, that is, beliefs must be consistent with equilibrium play. Note that rational-expectations beliefs trivially satisfy individual rationality.

We now turn to conditions on beliefs off the equilibrium path. Let $c^d_\tau \neq c^*_\tau$ denote a deviation contract by firm τ, with associated (off-path) belief $q^d_\tau := q(c^d_\tau, C^*_{-\tau})$. In general, there are several off-path beliefs q^d_τ that sustain the same equilibrium outcomes. In addition, some particular off-path beliefs may eliminate all equilibrium candidates. Requiring off-path beliefs to be individually rational typically does not substantially reduce the set of admissible off-path beliefs. To further restrict this set, we consider the following definition.

Definition 3 (Perfectly-Competitive Beliefs). *For a given equilibrium set of contracts C^* and on-path equilibrium beliefs $q^*_\tau := q(c^*_\tau, C^*_{-\tau})$, we say that off-path belief $q^d_\tau = q(c^d_\tau, C^*_{-\tau})$ is **perfectly competitive** if $q^d_\tau = q^*_\tau$.*

In words, an off-path belief q^d_τ is perfectly competitive if it is equal to the respective on-path belief, q^*_τ. Under perfectly-competitive beliefs, firms take workers' beliefs as given; they cannot use contract choice c^d_τ to manipulate workers' beliefs.

We then impose the following condition on off-path beliefs.

Condition 4 (Off-Equilibrium-Path Beliefs). *Let C^* denote an equilibrium set of contracts and $q^*_\tau := q(c^*_\tau, C^*_{-\tau})$ for all $\tau \in [0, F]$. For any deviation $c^d_\tau \neq c^*_\tau$ by firm τ, its associated belief q^d_τ must be individually rational and, whenever possible, perfectly competitive.*

In the spirit of perfect competition, Condition 4 implies that firms take q^*_τ as given: as long as contracts are compatible with the workers' participation constraints, firms cannot manipulate the proportion of each type of worker that applies to them. We need the "whenever possible" qualifier because, in some cases, setting $q^d_\tau = q^*_\tau$ will violate the individual rationality condition. This will happen if and only if, under belief $q^d_\tau = q^*_\tau$, we have $U_i(c^d_\tau, C^*_{-\tau}) < U^*_i(C^*_{-\tau})$ for either $i = b$ or $i = r$ (but not both). In such cases, Condition 4 requires only individual rationality, and beliefs must change after such a deviation. In what follows, if Condition 4 does not uniquely determine an off-path belief after deviation c_τ, whenever possible, we select an individually-rational belief q^d_τ such that deviation c^d_τ is not profitable.

We augment Definition 1 of equilibrium by further requiring Conditions 3 and 4 to hold. We denote an equilibrium by (C^*, μ, q), where μ and q are shortcuts for equilibrium residual labor supply and belief functions. We say that an equilibrium (C^*, μ, q) displays *full segregation* if $q(c^*_\tau, C^*_{-\tau}) \in \{0,1\}$ for all $\tau \in [0,F]$. That is, in a full segregation equilibrium, there are no firms with a mixed workforce. The following result shows that there is no full segregation equilibrium when firms cannot offer discriminatory contracts.

Lemma 1 (**No Full Segregation without Overt Discrimination**). *If overt discrimination is not possible, there is no equilibrium with full segregation.*

Intuitively, full segregation cannot happen because Blue workers would all apply to firms with only Red workers. This result implies that the full segregation equilibrium in Proposition 3 requires overt discrimination by firms.

We now state our main result.

Proposition 4 (**Partial Segregation Equilibrium**). *Let (C^*, μ, q) denote a tight-labor-market equilibrium under no overt discrimination. The equilibrium displays **partial segregation**:*

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

2. *A fraction $1 - s^*$ of "Red firms" hire one unit mass of Red workers for wages (w^*_{lr}, w^*_{hr}).*

Mixed firms pay higher entry-level wages than Red firms: $w^*_{lm} > w^*_{lr}$.

This proposition implies that, in equilibrium, firms differentiate themselves by choosing one of two contracts. "Mixed" firms employ both 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. These firms are large and pay high entry-level wages, so we refer to them as *high-wage firms*. Blue workers strictly prefer to work in such firms, because they have a higher probability of promotion than Red workers. "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. These firms are smaller and pay low entry-level wages, so we refer to them as *low-wage firms*. While such firms cannot discriminate in hiring, they still end up with a fully segregated workforce because only Red

workers find it optimal to apply to them. Red workers are indifferent between Mixed and Red firms.

The equilibrium is unequal in several ways, as the next corollary shows.

Corollary 1 (**Blue-Red Inequality**). *In a tight-labor-market equilibrium:*

i. Red workers are less likely than Blue workers to work for Mixed firms.

ii. In Mixed firms, Blue workers have higher promotion probabilities and average wages than Red workers.

iii. Blue workers have higher utility than Red workers.

Corollary 1 shows that the Blue-Red inequality is driven by both sorting and unequal treatment within firms. These two channels are linked. Because Red workers anticipate lower promotion probabilities than Blue workers in large, high-wage firms, Red workers are less likely than Blue workers to apply to these firms. Such results are consistent with the evidence in Card, Cardoso, and Kline (2016), who study gender gaps in employment and earnings across firms and find evidence of "(...) *a sorting channel that arises if women are less likely to be employed at higher-wage firms, and a bargaining channel that arises if women obtain a smaller share of the surplus associated with their job*." Similarly, when investigating the recent widening of the gender earnings gap, Goldin et al. (2017) show that this widening "*is split between men's greater ability or preference to move to higher paying firms and positions and their better facility to advance within firms*" (p. 114). Our model jointly explains such facts and the emergence of high-wage and low-wage firms.

An equilibrium must have Mixed firms but may or may not have Red firms. The following result shows a necessary condition for Red firms to exist.

Proposition 5 (**Bias and Firm Heterogeneity**). *There exists* $\overline{\beta} > 0$ *such that if* $\beta < \overline{\beta}$*, then* $s^* = 1$ *(i.e., no Red firms in equilibrium).*

Proposition 5 illustrates the importance of bias for firm heterogeneity. Unless the bias is sufficiently strong, all firms are identical in equilibrium (i.e., there is a corner solution: $s^* = 1$). The intuition is as follows. Perfect competition with ex-ante identical firms requires all firms to have the same profit in equilibrium. In an equilibrium with both types of firms, firms must be indifferent between the two contracts. If the bias is zero, neither firms nor workers care about the proportions of Red and Blue applying to each firm, implying that all firms face the same maximization problem. The optimal contract is the unique solution

to this problem. Given the assumption of a tight labor market, that contract must be a risky career path. Thus, at $\beta = 0$, a Mixed firm's profit must be strictly higher than that of a Red firm (i.e., a safe career path contract). Continuity then implies that a Mixed firm's profit is strictly higher than a Red firm's profit for β sufficiently close to zero, implying that no firm chooses to be Red.

3.7 Equilibrium Characterization

In this subsection, we present and interpret the conditions that characterize an equilibrium.

Proposition 4 implies that an equilibrium is a set of two unique contracts $\{c_r^*, c_m^*\}$ and a belief function q, such that 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 $s^* \in (0,1]$. In the proof of Proposition 4, we show that in Mixed firms, only the participation constraints of the Red workers bind. The equilibrium must then satisfy seven conditions:

(i) Workers' lifetime utilities are

$$\begin{aligned} U_r^* &:= u(w_{lm}^*) + \tfrac{1-q^*\beta}{n_m^*}\theta u(w_{hm}^*) \\ U_b^* &:= u(w_{lm}^*) + \tfrac{1+(1-q^*)\beta}{n_m^*}\theta u(w_{hm}^*) \end{aligned} \tag{24}$$

where $q^* := q(c_m^*; C_{-\tau_m}^*)$.

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

$$\pi_m^* = \max_{n_m, w_{lm}, w_{hm}} R_h - w_{hm} + n_m(R_l - w_{lm}) \tag{25}$$

subject to

$$u(w_{lm}) + \frac{1-q^*\beta}{n_m}\theta u(w_{hm}) \geq U_r^*. \tag{26}$$

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

$$\pi_r^* := \max_{w_{lr}, w_{hr}} R_h - w_{hr} + R_l - w_{lr} \tag{27}$$

subject to

$$u(w_{lr}) + \theta u(w_{hr}) \geq U_r^*. \tag{28}$$

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

$$\begin{cases} \text{If } \pi_m^* > \pi_r^*, \text{ then } s^* = 1 \\ \text{If } \pi_m^* = \pi_r^*, \text{ then } s^* < 1 \end{cases} \tag{29}$$

(v) If $s^* < 1$, Red workers are indifferent between contracts c_m^* and c_r^*:

$$U_r^* = u(w_{lr}^*) + \theta u(w_{hr}^*). \tag{30}$$

(vi) Labor markets clear:

$$\begin{cases} s^* q^* n_m^* F = \alpha_b E \\ (1 - s^* + s^*(1-q^*) n_m^*) F = (1-\alpha_b) E \end{cases}. \tag{31}$$

(vii) There is no profitable deviation: $\nexists\ c_\tau^d$ such that $\mu_{b\tau}(c_\tau^d, C_{-\tau}^*) = n_\tau^d q(c_\tau^d, C_{-\tau}^*)$ and $\pi(c_\tau^d) > \pi_m^*$.

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

Lemma 2. *Condition (vii) holds if and only if*

$$\pi_m^* \geq \max_c \pi(c) \ \ s.t.\ U_b(c;1) \geq U_b^*, \tag{32}$$

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

Intuitively, potentially profitable deviations involve hiring only Blue workers, in which case we must have $q^d = 1$. Thus, the absence of a profitable deviation requires the equilibrium profit to be greater than the maximum profit subject to hiring only Blue workers.

Similar to the benchmark case, three intuitive marginal conditions must hold in equilibrium. First, for Red firms we have:

$$u'(w_{lr}^*) = \theta u'(w_{hr}^*). \tag{33}$$

This condition implies that in Red firms, wages increase upon promotion, and that marginal utilities are equalized across jobs, from the perspective of Red workers—the only workers employed by Red firms. While Red firms choose wages efficiently given their career structures, the mere existence of these firms indicates resource misallocation because, in the benchmark model, 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, biased promotions in some firms distort economy-wide career opportunities.

Similarly, for Mixed firms we have

$$u'(w_{lm}^*) = \theta(1 - q^*\beta) u'(w_{hm}^*). \tag{34}$$

As before, wages increase upon promotion.[22] As in (10), marginal utilities are not equalized across jobs. 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.

Condition (34) implies that the equilibrium is inefficient. All else equal, in Mixed firms, this inefficiency manifests as a smaller promotion premium than in the benchmark model. The "all else equal" condition is a crucial qualifier. An increase in β also generates non-obvious general equilibrium effects. Specifically, a higher β lowers the outside options of Red workers, and hence their equilibrium utility U_r^*. What is less immediate is how firms respond. Facing lower worker bargaining power, firms can reduce the overall cost of contracts in different ways: they may lower wages (w_{lm}^* and w_{hm}^*), or expand hiring by increasing n_m^*, thereby reducing promotion probabilities. As a result, in Mixed firms, the relationship between U_r^* and wages is, in principle, ambiguous—unlike in the case of Red firms or in standard spot labor market models, where wages track utility more directly.

The following condition establishes a link between entry-level wages in Mixed firms and the Red workers' outside utility:

$$R_l = w_{lm}^* + \frac{U_r^* - u(w_{lm}^*)}{u'(w_{lm}^*)}. \tag{35}$$

This condition is equivalent to (19), and thus 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^*. Since equilibrium existence implies a unique solution to (35), w_{lm}^* increases when U_r^* decreases. If the equilibrium utility of Red workers decreases with β, as expected, entry-level wages in Mixed firms must *increase* with the bias. That is, higher promotion bias can increase the wage differential between high-wage firms and low-wage firms.

3.8 Effect of Bias on Equilibrium Outcomes

Since the equilibrium is characterized by a system of several nonlinear equations, closed-form comparative statics typically cannot be obtained without imposing strict functional form and parameter restrictions. We thus illustrate the equilibrium relationships using the same parametrization as in Figure 1. We focus on the case where $\beta \geq \overline{\beta}$, so that both Mixed and Red firms may coexist in equilibrium.[23]

[22] We note that $\theta(1 - q^*\beta) > 1$ in a tight-labor-market equilibrium.

[23] If $\beta < \overline{\beta}$, we have a corner solution where $s^* = 1$ (i.e., all firms are Mixed). In terms of predictions, this equilibrium is quite similar to the benchmark case with $\beta = 0$, although it also creates substantial inequality

Proposition 5 suggests that firm heterogeneity is related to the size of the bias. Thus, we first focus on the effect of the bias on the differences between Red and Mixed firms. Figure 2 presents equilibrium wages in Red and Mixed firms, as well as the size (i.e., span of control) of Mixed firms, n_m^*, as functions of β.

Panel A of Figure 2 shows wages in Red firms. As the bias increases, the outside option of Red workers, U_r^*, deteriorates, thereby lowering their bargaining power. Consequently, both entry-level and top wages decline with β.

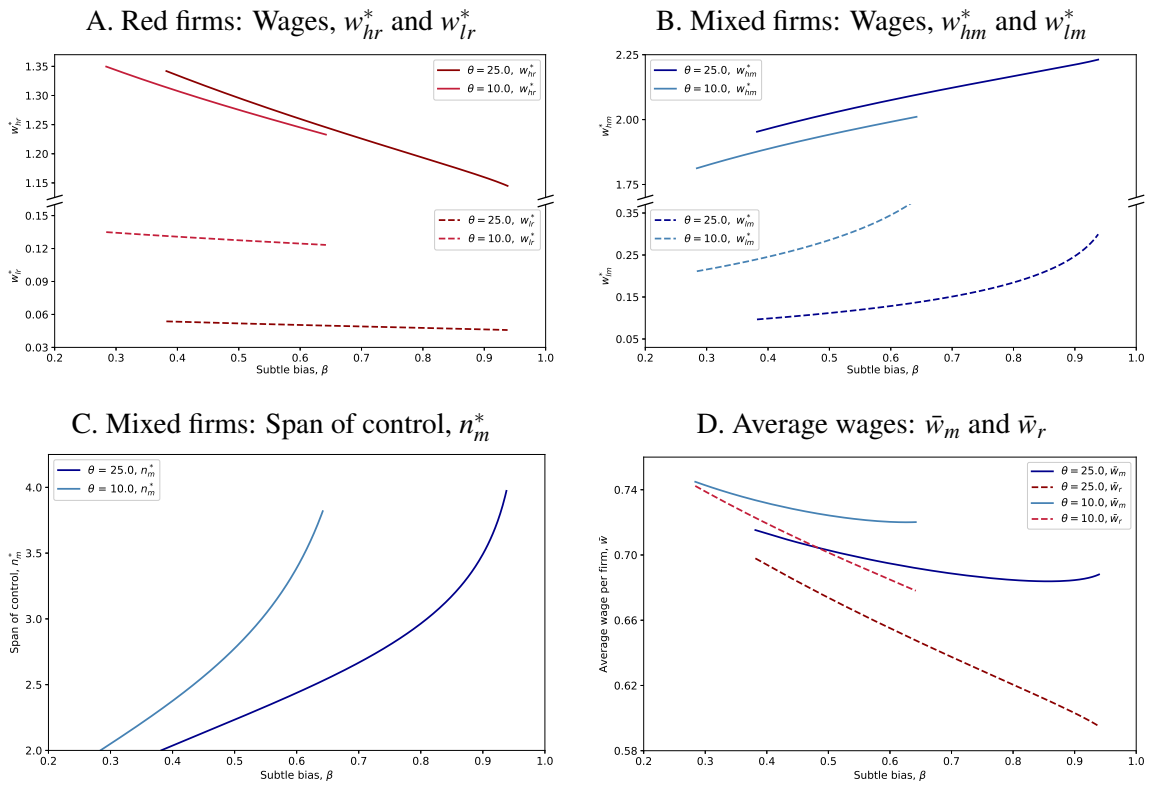


Figure 2: Optimal contracts and average wages in Red and Mixed firms under subtle bias.

Note: Panel A presents equilibrium wages in Red firms for the top job (w_{hr}^*, solid Red lines) and entry-level job (w_{lr}^*, dashed Red lines). Panel B shows corresponding wages in Mixed firms (w_{hm}^*, solid Blue lines; w_{lm}^*, dashed Blue lines). Panel C presents the equilibrium span of control in Mixed firms (n_m^*). Panel D shows average per-firm wages in Mixed firms ($\bar{w}_m$, solid Blue lines) and Red firms ($\bar{w}_r$, dashed Red lines). All panels plot outcomes as functions of subtle bias β, with $R_l = 0.75$ and $E/F = 2.0$.

between Blue and Red agents. We offer a full analysis of this equilibrium in the Internet Appendix.

Panel B of Figure 2 shows that Mixed firms pay more than Red firms at both levels. Proposition 4 shows that $w^*_{lm} > w^*_{lr}$. Here, we find that the same is true for top-level wages: $w^*_{hm} > w^*_{hr}$. Thus, Mixed firms are typically high-wage firms at all hierarchical levels.

Panel B also shows that in Mixed firms, both the entry-level wage and the top wage *increase* with the bias, which may seem counterintuitive. As firms become more biased against Red workers, the equilibrium utility U^*_r falls, which, from (35), implies that entry-level wages increase. At the same time, as workers become less powerful, the firms respond by expanding hiring to extract surplus from more workers. Panel C of Figure 2 indeed shows that firm size, or the *span of control*, n^*_m, increases with the bias, lowering the probability of promotion. To compensate for the lower promotion probability, Mixed firms increase wages.

Finally, Panel D of Figure 2 presents the average per-firm wages in Mixed and Red firms, computed as

$$\bar{w}_m = \frac{n^*_m w^*_{lm} + w^*_{hm}}{n^*_m + 1} \text{ and } \bar{w}_r = \frac{w^*_{lr} + w^*_{hr}}{2}. \tag{36}$$

We see that, also on average, Mixed firms are "high-wage firms" and Red firms are "low-wage firms."[24]

This example illustrates how subtle discrimination may lead to firm heterogeneity and partial segregation in equilibrium: large firms with risky promotion paths attract a mixed but Blue-dominated workforce and pay more on average, while small firms with seniority-based promotions attract only Red workers and pay less.

[24]Note that the fact that the Mixed firms have higher wages in both jobs does not imply that they also have higher average wages. In particular, if the span of control n^*_m is very large, the average wage $\bar{w}_m$ may fall below $\bar{w}_r$.

A. Share of Blue workers in Mixed firms, q^* B. Share of Mixed firms in the economy, s^*

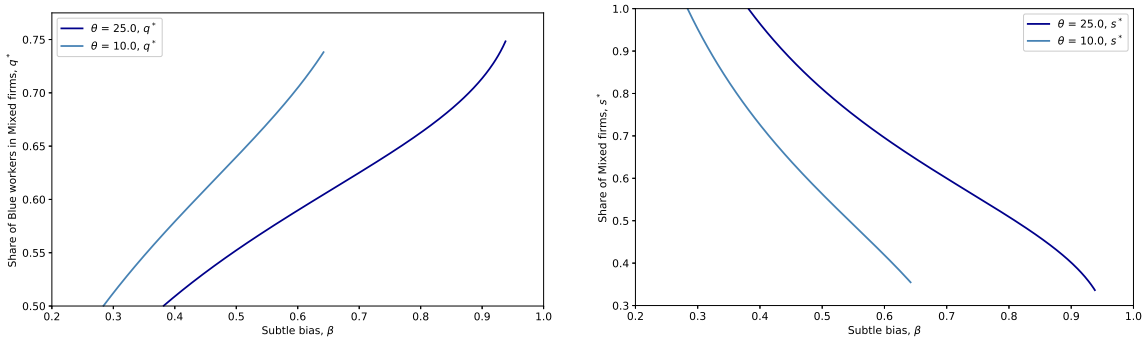


Figure 3: Share of Blue workers in Mixed firms and share of Blue firms in the economy.

Note: Panel A presents the equilibrium share of Blue workers in Mixed firms, q^*. Panel B shows the equilibrium share of Mixed firms in the economy, s^*. All panels plot outcomes as functions of subtle bias β, with $R_l = 0.75$ and $E/F = 2.0$.

Figure 3 presents the worker composition of Mixed firms (Panel A) and the firm type composition of the economy (Panel B) as a function of subtle bias. Panel A shows that as the bias increases, the high-wage Mixed firms become increasingly dominated by Blue workers. Panel B shows that the share of Mixed firms in the economy decreases with the bias. Thus, as bias rises, high-wage firms become less common but (inefficiently) larger, hiring many favored agents and a minority of unfavored agents. These firms offer risky career paths that lead to high wages upon promotion. By contrast, as the bias increases, the fraction of smaller, bureaucratic, low-wage firms increases in the economy. These firms specialize in hiring unfavored workers and offer predictable careers with lower wages at all corporate levels.

The Red workers who are promoted in high-wage firms are paid substantially more than those promoted in low-wage firms. However, the average wage of Red workers in Mixed firms is lower than that of Blue workers because of a lower promotion probability. All in all, Red workers are less likely to work in high-wage firms and, conditional on being employed by high-wage firms, have lower average earnings than Blue workers.

The careers in high-wage firms are unambiguously better than those in low-wage firms, in the sense that they yield higher expected utility under unbiased promotion rules. Favored workers therefore strictly prefer high-wage risky firms, while unfavored workers are indifferent between high-wage or low-wage firms. In practice, Blue workers disproportionately pursue high-wage, high-risk careers, while Red workers apply less frequently to these

positions and often opt for safer, lower-paying alternatives. This sorting pattern can be misinterpreted as reflecting inherent preference differences, suggesting that Red workers are more risk-averse or less financially motivated.

3.9 Welfare

The higher average wage for Blue workers is not a compensating differential for their higher promotion risk. In fact, Blue workers' lifetime utility is higher than that of Red workers, as shown in Figure 4. In other words, Blue workers enjoy rents simply for having the Blue label.

As β rises, Red workers become less powerful in Mixed firms, tightening their participation constraints. Because Mixed firms choose contracts that bind Red workers' participation constraints, weaker outside options for Red workers lead to less favorable terms for all workers. Although Blue workers are favored by the bias, as β increases, the fraction of Blue in Mixed firms, q^*, also increases, as well as firm size, n_m^*. The latter effect outweighs the positive bias effect, implying that Blue workers' utility also decreases with β. For sufficiently high β, both Blue and Red workers are worse off than in the egalitarian benchmark with no discrimination (shown as horizontal lines in Figure 4).

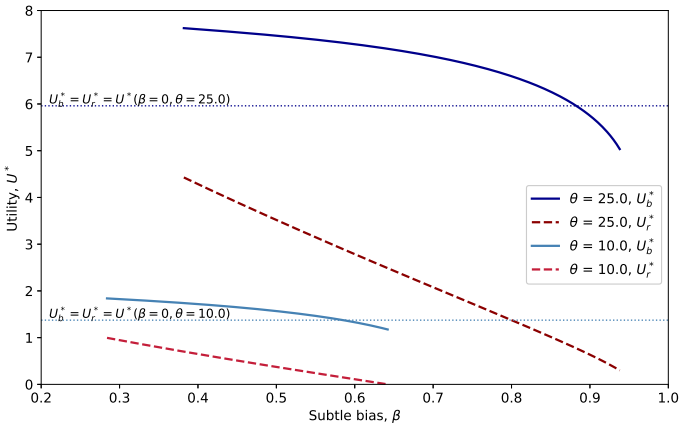


Figure 4: Utilities of Blue and Red workers, U_b^* and U_r^*.

Note: This figure presents the equilibrium utilities of Blue workers (solid blue lines) and Red workers (dashed red 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 5 shows that *profits increase with the bias*.[25] A larger bias benefits firms by reducing workers' bargaining power. Thus, firms may prefer the (inefficient) "biased equilibrium" to the (efficient) no-discrimination equilibrium: when β is sufficiently high, equilibrium profits exceed those in the no-discrimination case, as indicated by horizontal lines in Figure 5. Intuitively, the subtle bias acts as an implicit collusion device in the labor market: by depressing the outside options of Red workers, it allows firms to capture a larger share of a shrinking surplus. Perhaps unexpectedly, the model can rationalize postmodern 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 with perfect competition and rational expectations.

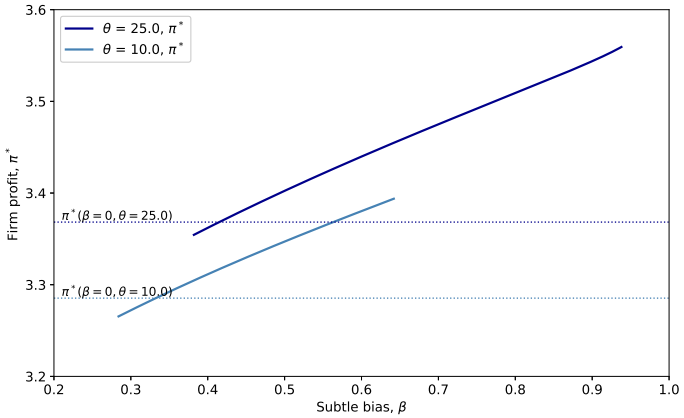


Figure 5: Firm profit, π^*.

Note: This figure presents the equilibrium profit of Mixed and Red firms, π^*, as a function of subtle bias, β, for $R_l = 0.75$, $R_h = 4.0$, $E/F = 2.0$. The horizontal dotted lines represent the equilibrium profit of Mixed firms, π^*, in the benchmark case, $\beta = 0$, for $\theta = 25.0$ and $\theta = 10.0$.

4 Conclusion

This paper presents a model of biased promotions where firm size, wages, and internal labor markets are endogenously determined in a labor market equilibrium. By providing a theoretical framework for understanding the consequences of biased promotions, our paper contributes to the broader literature on employer discrimination, firm performance, and inequality.

[25] In Figure 5, we set $R_h = 4.0$. Parameter R_h does not affect equilibrium outcomes and thus can assume any arbitrary positive level to achieve strictly positive profits.

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. In contrast, 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.

The framework is flexible and allows for several extensions, as illustrated in the Internet Appendix. In one such extension, we consider the case in which firms can make costly investments to correct their biases. We show that the equilibrium exhibits partial segregation among three types of firms: large Mixed firms, large Red firms, and small Red firms. Inequality and discrimination persist in this equilibrium, unless the cost of committing to treat both types equally goes to zero. In another extension, we show that the model predictions are robust to the unrestricted entry of 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.

A Appendix

Proof of Proposition 1. We assume that at least one tight-labor-market equilibrium exists. 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:

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

Because $E > F$, in any equilibrium, firms must offer $n^* > 1$, thus we set $\gamma^* = 0$. With $\gamma^* = 0$, 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." Both types must have the same equilibrium profit. Suppose a worker of type i has higher equilibrium utility than type $-i$. Then, a firm that hires type i workers can deviate and offer a similar contract with slightly lower wages to type $-i$ workers, who would accept. We conclude that both worker types must have the same utility in equilibrium. Thus, all firms face the same participation constraint. Proposition 1 then implies that both types of firms must offer identical contracts, except for the exclusion of one worker type. Because of full segregation, all workers in a firm face the same probability of promotion, which is $1/n^*$. Thus, they must be offered the unique contract described in Proposition 1. □

Proof of Lemma 1. Suppose C^* is a tight-labor market 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 per firm must be unique, n^*. $n^* = 1$ cannot be a tight-labor-market equilibrium because $E > F$. Therefore, all firms must offer $n^* > 1$. In this case, if firms are segregated, Blue workers would prefer to direct their search to firms where $q(c_\tau^*) = 0$. Thus, these firms must select some Blue workers, which is a contradiction. □

Proof of Proposition 4. Lemma 1 implies $\exists \tau \in [0,1]$ such that $q_\tau^* := q(c_\tau^*, C_{-\tau}^*) \in (0,1) \Rightarrow$ A Mixed firm must exist.[26]

Suppose first that this firm is such that $n_\tau^* = 1$ (i.e., a safe career path). Then, both types of workers have the same utility under c_τ^*. Because $E > F$, under a tight labor market there must be at least another firm such that $n_{\tau'}^* > 1$. If firm τ' is either Mixed or all Red, either the Red worker is worse off under $c_{\tau'}^*$ than c_τ^*, or the Blue worker would be better off under $c_{\tau'}^*$ than c_τ^*. In either case, c_τ^* and $c_{\tau'}^*$ cannot both be equilibrium contracts. If firm τ' is instead all Blue, we must have that the Blue worker is indifferent between c_τ^* and $c_{\tau'}^*$, and

[26] As discussed in Footnote 20, we assume that the equilibrium is such that $(1-q_\tau^*)\beta \leq n_\tau^* - 1$. If $(1-q_\tau^*)\beta > n_\tau^* - 1$, the equilibrium probabilities of promotion must be instead $p_b(n_\tau^*, q_\tau^*) = 1$ and $p_r(n_\tau^*, q_\tau^*) = \frac{1-n_\tau^* q_\tau^*}{n_\tau^* - n_\tau^* q_\tau^*}$. The analysis of this case is similar, and thus omitted for brevity.

also that $\pi(c^*_\tau) = \pi(c_{\tau'})$, otherwise the firm with the lowest profit could deviate and mimic the one with the highest profit. But then both firms face the same maximization problem as in (A.1), of which $n^*_{\tau'} > n^*_\tau = 1$ cannot both be optimal solutions. We conclude that the Mixed firm must have $n^*_\tau > 1$.

We now show that all Mixed firms offer the same contract, where Red's participation constraint binds and Blue's participation constraint is slack. Consider a Mixed firm such that $n^*_\tau > 1$. Conditions 2-4 and profit maximization imply

$$\pi(c^*_\tau) = \max_{w_l, w_h, n} R_h - w_h + n(R_l - w_l) \tag{A.11}$$

subject to

$$u(w_l) + \frac{1 - q^*_\tau \beta}{n} \theta u(w_h) \geq U^*_r(C^*_{-\tau}) \tag{A.12}$$

$$u(w_l) + \frac{1 + (1 - q^*_\tau)\beta}{n} \theta u(w_h) \geq U^*_b(C^*_{-\tau}) \tag{A.13}$$

$$n \geq 1. \tag{A.14}$$

One of (A.12) and (A.13) must bind, otherwise, firm τ can always increase n to increase profit. Suppose the Blue constraint (A.13) binds. Consider first a deviation contract c^d_τ that maximizes profit while ignoring the Blue constraint. Let $q^d_\tau := q(c^d_\tau, C^*_{-\tau})$ denote the belief after this deviation. Because this is an "optimal deviation," the Red constraint must bind (otherwise the firm can increase n^d_τ, which relaxes (A.14) and increases profit), implying $U_r(c^d_\tau; q^d_\tau) = U^*_r(C^*_{-\tau})$, where $U_i(c; q)$ denotes the utility of a type-i agent that accepts contract c under belief q.

Note that in equilibrium, the Blue worker has strictly higher utility than the Red worker:

$$U^*_b(C^*_{-\tau}) = U_b(c^*_\tau; q^*_\tau) = U_r(c^*_\tau; q^*_\tau) + \frac{\beta}{n^*_\tau} \theta u(w^*_{h\tau}) > U^*_r(C^*_{-\tau}). \tag{A.15}$$

If c^d_τ does not violate the Blue constraint, Condition 4 implies that $q^d_\tau = q^*_\tau$, and thus profit maximization implies that $\pi(c^d_\tau) \leq \pi(c^*_\tau) \Rightarrow$ the deviation is not profitable. Thus, suppose instead that c^d_τ violates the Blue constraint. Condition 4 now implies that $q^d_\tau = 0$, which means only Red workers are matched to the deviating firm.

Suppose first that (A.14) also binds (in addition to the Red constraint), i.e., $n^d_\tau = 1$. Because $U_r(c^d_\tau; q^d_\tau) = U^*_r(C^*_{-\tau})$, Blue workers strictly prefer some contract in $C^*_{-\tau}$ to c^d_τ, which confirms that no Blue worker applies to contract c^d_τ. Thus, c^d_τ is a profitable deviation for the firm $\Rightarrow$ the Blue constraint (A.13) cannot bind in a equilibrium where firm τ hires both Red and Blue workers.

Suppose now that, under the optimal deviation, $n_\tau^d > 1$. We then have

$$U_b(c_\tau^d; q_\tau^d) = U_r(c_\tau^d; q_\tau^d) + \beta\theta \frac{u(w_{h\tau}^d)}{n_\tau^d} = U_r^*(C_{-\tau}^*) + \beta \frac{R_l - w_{l\tau}^d}{(1 - q_\tau^d \beta)} u'(w_{l\tau}^d), \tag{A.16}$$

where the last step follows from the first-order conditions when the Red constraint binds. When the Red constraint binds, $w_{l\tau}^d$ is independent of q_τ^d. Thus, $U_b(c^d; q)$ increases with q, which implies $U_b(c_\tau^d; 0) < U_b^*(C_{-\tau}^*)$, which confirms that no Blue worker applies to contract c_τ^d. Thus, c_τ^d is a profitable deviation for the firm $\Rightarrow$ the Blue constraint (A.13) cannot bind in an equilibrium where firm τ hires both Red and Blue workers.

Since a Mixed firm must have more than a unit mass of workers and the Red constraint must bind, there is only one contract that maximizes profit for Mixed firms, implying that all Mixed firms are identical. We denote the optimal contract for Mixed firms by $c_m^* = (w_{lm}^*, w_{hm}^*, n_m^*)$

Suppose now there exists an optimal contract $c_r^* = (w_{lr}^*, w_{hr}^*, n_r^*)$ such that $n_r^* = 1$. Because Mixed firms attract Red workers, $U_r(c_m^*; q_m^*) \geq U_r(c_r^*; q_r^*) = U_b(c_r^*; q_r^*)$, where the last step follows because under c_r^*, the probability of promotion is one. Since $U_b(c_m^*; q_m^*) > U_r(c_m^*; q_m^*)$, we conclude that c_r^* may exist only if it is a fully Red firm, i.e., $q_r^* = 0$.

We now show that the entry-level wages in Red firms are lower than the entry-level wages in Mixed firms (i.e. $w_{lm}^* > w_{lr}^*$). The equilibrium contract in Red firms must satisfy the marginal condition $u'(w_{lr}^*) = \theta u'(w_{hr}^*)$, and the equilibrium contract in Mixed firms must be such that $u'(w_{lm}^*) = \theta(1 - q^*\beta)u'(w_{hm}^*)$. We rearrange the marginal conditions:

$$\frac{u'(w_{lr}^*)}{u'(w_{lm}^*)} = \frac{u'(w_{hr}^*)}{(1 - q^*\beta)u'(w_{hm}^*)}.$$

Suppose $w_{lr}^* \geq w_{lm}^*$. Then, we have

$$\frac{u'(w_{lr}^*)}{u'(w_{lm}^*)} = \frac{u'(w_{hr}^*)}{(1 - q^*\beta)u'(w_{hm}^*)} \leq 1 \Rightarrow u'(w_{hr}^*) < u'(w_{hm}^*) \Rightarrow w_{hr}^* > w_{hm}^*.$$

Because in equilibrium a Red worker must be indifferent between the two contracts,

$$u(w_{lm}^*) + p_r \theta u(w_{hm}^*) = u(w_{lr}^*) + \theta u(w_{hr}^*)$$

or, rearranging,

$$u(w_{lm}^*) - u(w_{lr}^*) = \theta[u(w_{hr}^*) - p_r u(w_{hm}^*)], \tag{A.17}$$

where $p_r = \frac{1 - q^*\beta}{n_m^*} < 1$. If $w_{lr}^* \geq w_{lm}^*$ and $w_{hr}^* > w_{hm}^*$, the left-hand side of (A.17) is non-positive and the right-hand side is strictly positive, which is a contradiction. Thus, we must have $w_{lr}^* < w_{lm}^*$. □

Proof of Corollary 1: Blue workers work in Mixed firms with probability one. Red workers work in Mixed firms with probability

$$\frac{s^* n_m^* (1-q^*)}{s^* n_m^* (1-q^*) + (1-s^*)} \leq 1.$$

Blue workers' probability of promotion is

$$p_b(n_m^*, q^*) = \frac{1+(1-q^*)\beta}{n_m^*} > \frac{1-q^*\beta}{n_m^*} = p_r(n_m^*, q^*).$$

Expected wages are $w_{lm}^* + p_i(n_m^*, q^*) w_{hm}^*$. □

Proof of Proposition 5. With $\beta = 0$, the maximization problem of Mixed firms is:

$$\max_{w_{lm}, w_{hm}, n_m} R_h - w_{hm} + n_m (R_l - w_{lm}) \tag{A.18}$$

subject to

$$\begin{cases} u(w_{lm}) + \frac{1}{n_m} \theta u(w_{hm}) \geq U_r^* \\ n_m \geq 1 \end{cases} \tag{A.19}$$

For a given U_r^*, the resulting optimal contract is unique (see Proposition 1) and in order to have a tight labor market equilibrium the parameters are such that $n_m^* > 1$, it then follows that the profit from offering a career path is higher than the profit from offering a safe career path, and therefore there is only one type of firms (i.e., Mixed firms) in equilibrium. □

Proof of Lemma 2. In the proof of Proposition 4, we show that the Red constraint must bind in equilibrium. Therefore, a 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, no one will be assigned 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. □

References

ARELLANO-BOVER, J., N. BIANCHI, S. LATTANZIO, AND M. PARADISI (2024): "One Cohort at a Time: Young Workers' Outcomes and the Declining Gender Pay Gap," NBER Working Paper 32612, National Bureau of Economic Research.

ATHEY, S., C. AVERY, AND P. ZEMSKY (2000): "Mentoring and diversity," *American Economic Review*, 90, 765–786.

AURIOL, E., G. FRIEBEL, AND F. VON BIEBERSTEIN (2016): "The firm as the locus of social comparisons: Standard promotion practices versus up-or-out," *Journal of Economic Behavior & Organization*, 121, 41–59.

AURIOL, E. AND R. RENAULT (2008): "Status and incentives," *The RAND Journal of Economics*, 39, 305–326.

AXELSON, U. AND P. BOND (2015): "Wall Street Occupations," *Journal of Finance*, 70, 1949–1996.

AZMAT, G., V. CUÑAT, AND E. HENRY (2025): "Gender promotion gaps and career aspirations," *Management Science*, 71, 2127–2141.

BABCOCK, L., M. P. RECALDE, L. VESTERLUND, AND L. WEINGART (2017): "Gender differences in accepting and receiving requests for tasks with low promotability," *American Economic Review*, 107, 714–747.

BAKER, G., M. GIBBS, AND B. HOLMSTROM (1994): "The Internal Economics of the Firm: Evidence from Personnel Data," *Quarterly Journal of Economics*, 109, 881–919.

BAKER, G. P., M. C. JENSEN, AND K. J. MURPHY (1988): "Compensation and incentives: Practice vs. theory," *The Journal of Finance*, 43, 593–616.

BARTOŠ, V., M. BAUER, J. CHYTILOVÁ, AND F. MATĚJKA (2016): "Attention discrimination: Theory and field experiments with monitoring information acquisition," *American Economic Review*, 106, 1437–75.

BECKER, G. S., K. M. MURPHY, AND I. WERNING (2005): "The equilibrium distribution of income and the market for status," *Journal of Political Economy*, 113, 282–310.

BENNEDSEN, M., E. SIMINTZI, M. TSOUTSOURA, AND D. WOLFENZON (2022): "Do Firms Respond to Gender Pay Gap Transparency?" *The Journal of Finance*, 77, 2051–2091.

BENSON, A., D. LI, AND K. SHUE (2019): "Promotions and the Peter principle," *The Quarterly Journal of Economics*, 134, 2085–2134.

——— (2024): ""Potential" and the gender promotion gap," *American Economic Review*, forthcoming.

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

BERTRAND, M. AND K. F. HALLOCK (2001): "The gender gap in top corporate jobs," *Ilr Review*, 55, 3–21.

BIDWELL, M. J., S. WON, R. BARBULESCU, AND E. MOLLICK (2015): "I Used to Work at Goldman Sachs! How Firms Benefit from Organizational Status in the Market for Human Capital," *Strategic Management Journal*, 36, 1164–1173.

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

——— (2024): "Gender promotion gaps in knowledge work: The role of task assignment in teams," .

BOHREN, J. A., A. IMAS, AND M. ROSENBERG (2019): "The dynamics of discrimination: Theory and evidence," *American Economic Review*, 109, 3395–3436.

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.

CARD, D., A. R. CARDOSO, J. HEINING, AND P. KLINE (2018): "Firms and labor market inequality: Evidence and some theory," *Journal of Labor Economics*, 36, S13–S70.

CARD, D., A. R. CARDOSO, AND P. KLINE (2016): "Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women," *The Quarterly Journal of Economics*, 131, 633–686.

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

CULLEN, Z. AND R. PEREZ-TRUGLIA (2023): "The old boys' club: Schmoozing and the gender gap," *American Economic Review*, 113, 1703–1740.

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

DAVIES, S., E. VAN WESEP, AND B. WATERS (2024): "On the Magnification of Small Biases in Hiring," *The Journal of Finance*, 79, 3623–3673.

DOERINGER, P. B. AND M. J. PIORE (1971): *Internal Labor Markets and Manpower Analysis*, Lexington, MA: Heath Lexington Books.

EWENS, M. AND X. GIROUD (2024): "Corporate Hierarchy," Working Paper.

FAHN, M. AND N. KLEIN (2025): "The Dynamics of Exploiting Overconfident Workers," *Journal of Labor Economics*, forthcoming.

FANG, H. AND A. MORO (2011): "Theories of statistical discrimination and affirmative action: A survey," in *Handbook of Social Economics*, ed. by J. Benhabib, A. Bisin, and M. O. Jackson, North-Holland, vol. 1, chap. 5, 133–200.

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

——— (2025): "Polarization, purpose and profit," *Journal of Financial Economics*, 172, 104147.

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

GIBBONS, R. AND M. WALDMAN (1999): "A Theory of Wage and Promotion Dynamics Inside Firms," *Quarterly Journal of Economics*, 114, 1321–1358.

GOLDIN, C., S. P. KERR, C. OLIVETTI, AND E. BARTH (2017): "The Expanding Gender Earnings Gap: Evidence from the LEHD-2000 Census," *American Economic Review: Papers and Proceedings*, 107, 110–114.

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

HOPKINS, E. AND T. KORNIENKO (2004): "Running to Keep in the Same Place: Consumer Choice as a Game of Status," *American Economic Review*, 94, 1085–1107.

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.

HUBER, K., V. LINDENTHAL, AND F. WALDINGER (2021): "Discrimination, managers, and firm performance: Evidence from "aryanizations" in Nazi Germany," *Journal of Political Economy*, 129, 2455–2503.

HUITFELDT, I., A. R. KOSTØL, J. NIMCZIK, AND A. WEBER (2023): "Internal labor markets: A worker flow approach," *Journal of Econometrics*, 233, 661–688.

HUNEEUS, F., C. MILLER, C. NEILSON, AND S. D. ZIMMERMAN (2021): "Firm Sorting, College Major, and the Gender Earnings Gap," *Working Paper, Princeton University*.

KE, R., J. LI, AND M. POWELL (2018): "Managing Careers in Organizations," *Journal of Labor Economics*, 36, 197–252.

KLINE, P., E. K. ROSE, AND C. R. WALTERS (2022): "Systemic Discrimination Among Large U.S. Employers," *The Quarterly Journal of Economics*, 137, 1963–2036.

KREPS, D. M. AND J. A. SCHEINKMAN (1983): "Quantity precommitment and Bertrand competition yield Cournot outcomes," *The Bell Journal of Economics*, 326–337.

KUNZE, A. AND A. R. MILLER (2017): "Women helping women? Evidence from private sector data on workplace hierarchies," *Review of Economics and Statistics*, 99, 769–775.

LANG, K. AND J.-Y. K. LEHMANN (2012): "Racial discrimination in the labor market: Theory and empirics," *Journal of Economic Literature*, 50, 959–1006.

LANG, K., M. MANOVE, AND W. T. DICKENS (2005): "Racial discrimination in labor markets with posted wage offers," *American Economic Review*, 95, 1327–1340.

LAZEAR, E., K. SHAW, AND C. STANTON (2018): "Who gets hired? The importance of competition among applicants," *Journal of Labor Economics*, 36, S133–S181.

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

——— (1990): "Male-female wage differentials in job ladders," *Journal of Labor Economics*, 8, S106–S123.

MATSA, D. A. AND A. R. MILLER (2011): "Chipping away at the glass ceiling: Gender spillovers in corporate leadership," *American Economic Review*, 101, 635–639.

MINNI, V., K.-T. NGUYEN, AND H. SARSONS (2025): "Managers and the Cultural Transmission of Gender Norms," Working Paper.

ONUCHIC, P. (2023): "Recent contributions to theories of discrimination," Working paper, Oxford University.

ONUCHIC, P. AND D. RAY (2023): "Signaling and discrimination in collaborative projects," *American Economic Review*, 113, 210–252.

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

PRENDERGAST, C. (1993): "The role of promotion in inducing specific human capital acquisition," *The Quarterly Journal of Economics*, 108, 523–534.

RAY, D. AND A. ROBSON (2012): "Status, Intertemporal Choice, and Risk-Taking," *Econometrica*, 80, 1505–1531.

RONCHI, M. AND N. SMITH (2021): "Daddy's girl: Daughters, managerial decisions, and gender inequality,," Working paper, Bocconi University.

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.

SINISCALCHI, M. AND P. VERONESI (2021): "Self-image bias and lost talent," Working paper, NBER.

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

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

WRIGHT, R., P. KIRCHER, B. JULIEN, AND V. GUERRIERI (2021): "Directed search and competitive search equilibrium: A guided tour," *Journal of Economic Literature*, 59, 90–148.