Promotion signaling, discrimination, and positive discrimination policies

22 March 2018, 11:15-12:30, bld. 72, room 465

Oliver Gürtler (University of Cologne)

Abstract: This paper studies discrimination in a model in which promotions are used as signals of worker ability. The model can account for statistical and taste-based discrimination. In the short run, a positive discrimination policy is beneficial for workers in the middle of the ability distribution because these workers are promoted if and only if the policy is in place. Instead, workers of either high or low ability suffer from the policy. In the long run, the policy benefits all targeted workers. The model can explain empirical findings about the effects of a gender quota on the boards of Norwegian companies.

Promotion signaling, discrimination, and positive discrimination policies*

Marc Gürtler**, *Braunschweig Institute of Technology*

Oliver Gürtler***, *University of Cologne*

Abstract

This paper studies discrimination in a model in which promotions are used as signals of worker ability. The model can account for statistical and taste-based discrimination. In the short run, a positive discrimination policy is beneficial for workers in the middle of the ability distribution because these workers are promoted if and only if the policy is in place. Instead, workers of either high or low ability suffer from the policy. In the long run, the policy benefits all targeted workers. The model can explain empirical findings about the effects of a gender quota on the boards of Norwegian companies.

Keywords: promotion, signaling, discrimination, positive discrimination policy
JEL classification: D83, J24, J31, J78, M51

*We thank Jed DeVaro, Christian Ewerhart, Matthias Kräkel, Patrick Legros, Johannes Münster, Hideo Owan, Martin Ruckes, Patrick Schmitz, Anders Stenberg, and participants at the 41st Annual Conference of the Eastern Economic Association, the 18th Colloquium on Personnel Economics, the Annual 2015 Conference of the Verein für Socialpolitik, the 85th Annual Meetings of the Southern Economic Association, the Econ and Finance Workshop in Karlsruhe, the research seminar ULME in Ulm, the DICE research seminar in Düsseldorf, the DART workshop in Graz, and the research seminar at Linnaeus University for helpful comments.

**Marc Gürtler, Department of Economics, Braunschweig Institute of Technology, Abt-Jerusalem-Str. 7, D-38106 Braunschweig, Germany. Phone: +49-531-3912895; E-mail: marc.guertler@tu-bs.de

***Oliver Gürtler, Department of Economics, University of Cologne, Albertus Magnus Platz, D-50923 Cologne, Germany. Phone: +49-221-4701450; E-mail: oliver.guertler@uni-koeln.de

1 Introduction

Positive discrimination policies are implemented all over the world with the aim of promoting the opportunities of people who (are perceived to) suffer from discrimination. These policies can differ in their nature: some are intended to facilitate obtainment of higher education, while others forbid firms to discriminate against people based on their race, religion, sex, and so on, when making hiring or promotion decisions. In this paper, we consider positive discrimination policies that are aimed at improving the career prospects of people who are discriminated against by ordering firms to make it easier for those people to be promoted.[1] A recent example of such a policy is the European Commission's proposal to impose a 40-percent quota for female directors on the supervisory boards of publicly listed companies.[2] We show that these policies may be counterproductive in that they can actually hurt the people they are intended to benefit.

In the economic literature, it has long been recognized that promotions can serve as signals about worker ability.[3] The intuition is simple. A worker's current employer is likely to receive more comprehensive information about the worker's ability than external firms, i.e., learning is asymmetric. When performance is more sensitive to ability in high-level jobs rather than in low-level jobs, the current employer promotes the worker to a high-level job if and only if the employer believes that the worker's ability is sufficiently high. As a consequence, when external firms observe a worker's promotion, they upgrade their assessment of the worker's ability. In turn, they have a greater interest in hiring that worker, with the result that the worker receives more generous wage offers. Finally, because firms must pay higher wages to retain promoted workers, firms decide to promote inefficiently few workers; that is, promotion standards are inefficiently high. Recent empirical studies find results that are in line with predictions derived from the promotion-signaling model.[4]

We consider a promotion-signaling model with two periods. In the first period, firms

[1] Fryer and Loury (2005) argue that positive discrimination policies are highly controversial and therefore emphasize the importance of economic reasoning in the evaluation of these programs.

[2] http://ec.europa.eu/justice/newsroom/gender-equality/news/121114_en.htm#Press (last retrieved on October 12, 2017).

[3] See, e.g., Waldman (1984), Bernhardt (1995), Zábojník and Bernhardt (2001), Owan (2004), Ghosh and Waldman (2010), Zábojník (2012), Waldman (2013), Gürtler and Gürtler (2015), DeVaro and Kauhanen (2016), Shankar (2016), Waldman (2016), and Waldman and Zax (2016).

[4] See DeVaro and Waldman (2012), Bognanno and Melero (2016), and Cassidy et al. (2016).

hire workers and assign them to a low-level job. Each firm observes the employed workers' performances at the end of the first period and then decides which workers to promote to a high-level job and which to reassign to the low-level one. External firms observe job assignment decisions, and thereafter make wage offers to the workers to hire them. In a first step, we show that the model can be used to explain why some workers are discriminated against. In particular, the model can capture both (endogenous and exogenous) statistical discrimination and taste-based discrimination. Take endogenous statistical discrimination as an example. It is optimal for the first-period employer to promote a worker when first-period performance exceeds a certain performance standard. This performance standard need not be uniquely defined. The intuition for this result is the following: if external firms expect a worker's first-period employer to set a rather high promotion standard, they conclude that a promoted worker must be of exceptionally high ability. Therefore, they make very generous wage offers to promoted workers, which in turn makes it optimal for the first-period employer to reassign the worker to the low-level job unless his performance is exceptionally high. In other words, the external firms' expectation of a very high promotion standard may become self-fulfilling in the current model. The same is true for a relatively low promotion standard. It is then conceivable that there are two identical workers playing different equilibria with different promotion standards. In turn, only one of them (the worker with the lower promotion standard) may be promoted, whereas the other worker is reassigned to the low-level job and thus discriminated against.

In a second step, we introduce a positive discrimination policy into the model aimed at improving the career prospects of people who are discriminated against. In particular, we assume that a quota is introduced, requiring that workers from an originally disadvantaged group fill a certain share of the positions in the high-level job. We show that introducing the quota lowers the promotion standard for workers who were originally disadvantaged and, thus, favored by the policy and it increases the promotion standard for workers disadvantaged by the policy. We distinguish between the effects of the policy on workers who have already been hired when the policy is introduced, and which we term the policy's short-run effects, and the policy's long-run effects on workers who begin their career after the policy was introduced.

Consider first a worker in the first half of his career (i.e., the first period), who has already been hired when the policy is introduced and who is favored by the policy. For such a worker, the first-period wage does not change; the policy only affects the worker's second-period payoff. There are two effects. First, the worker is more likely to be promoted and to obtain a wage increase, which obviously benefits him or her. Second, the positive signal of promotion to the high-level job becomes weaker, whereas the negative signal of being reassigned to the low-level job becomes stronger. If the worker is promoted, external firms may question the worker's ability and may believe that the worker was promoted only because of the positive discrimination program. If, instead, the worker is not promoted, external firms believe that the worker's ability must be extremely low because he or she did not manage to become promoted in spite of the positive discrimination program. Because external firms interpret the job-assignment signal differently when a positive discrimination policy is in place, their wage offers differ as well. In particular, both the wage for a promoted worker and for a worker reassigned to the low-level job decreases, which clearly hurts workers. Summing up, the introduction of the positive discrimination policy leaves those workers worse off who have either very high ability, such that they would have been promoted even without the policy, or who have such low ability that they are not considered for promotion even when the policy is in place. In contrast, workers in the middle of the ability distribution benefit from the policy because those workers are promoted if and only if it is in place. It is possible that the negative effects on the workers of either low or high ability outweigh the positive effects on the workers of middle ability so that a worker's expected payoff may actually decrease.

Positive discrimination policies are sometimes criticized for devaluing the achievements of people who are intended to benefit from the policies, possibly leading to feelings of inferiority, self-doubt, and incompetence.[5] This argument is reminiscent of our finding that the positive signal of promotion is weaker when a positive discrimination policy is in place than when there is no such policy. As we show, the problem may be less that promoted workers question their own ability, but rather that external firms do, leading to less generous wage offers.

Consider now a worker who begins his working career only after a positive discrimination policy has been introduced. The policy affects such a worker's payoff in three additional ways. First, the change in the promotion standard alters the worker's expected second-period compensation, as just explained, which in turn affects the firms' willingness to pay for the worker in the first period. Second, by changing the number of workers from the different worker groups initially hired, firms can affect the promotion standards. For instance, if a

[5] See, for example, Andre et al. (1992).

firm hires many workers favored by the policy, but only few workers disadvantaged by the policy, the firm finds it relatively easy to satisfy the quota constraint, so that the promotion thresholds do not deviate strongly (or at all) from the standards derived in the model without the positive discrimination policy. In other instances, the quota constraint may be more difficult to satisfy, implying bigger changes in the promotion thresholds. As we will show, firms make hiring decisions such that optimal promotion standards are implemented in the sense that workers' expected average output is maximized. The corresponding change in workers' outputs again affects the firms' willingness to pay for workers' services. Third, due to the policy, hiring decisions for different workers are no longer independent. This implies that a worker may receive a total expected wage that differs from the worker's expected productivity, influencing the worker's payoff as well. As a consequence of these additional effects, we find that the targeted workers benefit from the policy regardless of their exact ability in the long run.

Positive discrimination policies are often argued to lead to mismatching, thereby reducing productivity and efficiency.[6] Our results demonstrate that this is not necessarily true. As mentioned before, firms react to the policy introduction by changing the number of workers emanating from different worker groups hired into the low-level job to cope optimally with the policy introduction. As a result, promotion standards are implemented such that the workers' expected average output is maximized, resulting in a more productive assignment of workers to jobs.

We also investigate the effect of the policy on workers who are disadvantaged by the policy and therefore less likely to be promoted. Extremely able workers who are promoted in spite of the rule may benefit from the policy because the positive signal of being promoted grows stronger.[7]

[6] See Andre et al. (1992).

[7] This finding is related to the observation by Fang and Norman (2006) that people who (are thought to) suffer from government-mandated discriminatory policies are at times economically more successful than people who benefit from these policies. Fang and Norman (2006) present the example of ethnic Chinese in Malaysia, who were discriminated against by governmental implementation of the "New Economic Policy", but who did surprisingly well economically afterwards. Fang and Norman (2006) explain this observation using a model with a public and a private sector. If a group of people is discriminated against in the public sector, this group faces higher incentives to succeed in the private sector, implying that the group may perform better economically.

We derive several empirical implications from the model. An important one is that the introduction of a positive discrimination policy lowers the wages that workers who are supposed to be advantaged by the policy earn later in their career, and increases wages for those workers who are disadvantaged by the policy when controlling for job level. For workers who began their careers before the policy was put in place, their starting wage is unaffected by the policy. This in turn implies that the wage increase upon promotion of these workers is altered by the positive discrimination policy. For workers who are advantaged by the policy, the post-promotion wage decreases, while the starting wage is fixed so that the wage increase upon promotion becomes lower. For workers who are disadvantaged by the policy, the effect is the opposite and the wage increase upon promotion gets higher. Bertrand et al. (2014) study the effects of the introduction of a quota on the boards of publicly traded Norwegian companies, which led to an increase in the share of women elected to these boards. They show that election to a company board entails a substantial financial reward for the elected worker. After the board quota was introduced, this reward fell for women (from 9.4% percent of annual earnings to 8%), whereas it increased substantially for men (from 4.6% to 10%), in line with the predictions of our model. Several studies find that positive discrimination policies are rather ineffective in diminishing average "wage gaps" between different groups of workers.[8] Our model shows that positive discrimination policies can have adverse wage effects for the targeted workers, at least in the later stages of their careers, providing a possible explanation for this observation.

The paper is organized as follows. In the next section, we present related literature. In Section 3, we turn to the basic model. The effect of positive discrimination policies on workers' wages is analyzed in Section 4. The empirical implications of the model are discussed in Section 5, while Section 6 presents the conclusions. The proofs of all lemmas and propositions are in the Appendix.

2 Related literature

As indicated before, our paper makes a contribution to the promotion-signaling literature. The promotion-signaling model was developed by Waldman (1984) and extended by Bernhardt (1995), Zábojník and Bernhardt (2001), Owan (2004), Ghosh and Waldman (2010), De-

[8] See, for example, Smith and Welch (1984), Leonard (1996), and Burger and Jafta (2012).

Varo and Waldman (2012), DeVaro et al. (2012), Zábojník (2012), Waldman (2013), Gürtler and Gürtler (2015), Cassidy et al. (2016), DeVaro and Kauhanen (2016), Shankar (2016), Waldman (2016), and Waldman and Zax (2016). The current paper extends this literature by demonstrating how different types of discrimination can occur in the promotion-signaling model. Furthermore, to our knowledge it is the first paper that examines the consequences of positive discrimination policies in that model.

The paper is related to the economic literature on discrimination. This literature can be divided into two different strands.[9] First, there is a body of work on taste-based discrimination that was originated by Becker (1957) and further developed by, e.g., Coate and Loury (1993a) and Black (1995). According to this literature, workers belonging to some group of people are discriminated against because firms incur some disutility when interacting with these workers. Second, there is a body of literature on statistical discrimination. In this literature, it is assumed that a worker's ability is not fully known to potential employers. Firms therefore use all the available information to estimate abilities. When there are differences between groups, these differences influence the ability assessments so that it makes sense for the firms to treat two workers who belong to two different groups, but are otherwise identical, differently. Differences between groups can either be imposed exogenously (e.g., Phelps 1972), or can emerge endogenously (e.g., Coate and Loury 1993b, Moro and Norman 2003, Fryer 2007). The intuition for the latter possibility is as follows. When firms hold pessimistic beliefs about the abilities of a certain group of people, they are unwilling to hire these people. In turn, members of this group of people have low incentives to enhance their abilities, thus confirming the firms' pessimistic beliefs.

Our paper contributes to this literature in two ways. First, we show how discrimination may arise in a promotion-signaling framework. Second, we examine the effects of positive discrimination policies on the payoffs of the affected workers. In many of the papers mentioned above, the effect of positive discrimination policies on the model outcome is analyzed. In contrast to our study, however, wages are often assumed to be exogenous. One paper that endogenizes wages is by Moro and Norman (2003), who study statistical discrimination in a general equilibrium model with two different types of jobs. As in our model, they show that positive discrimination policies can lower the wages of workers who are supposed

[9] See Fang and Moro (2011) as well as Lang and Lehmann (2012) for surveys of the economic literature on discrimination.

to benefit from the policy. This happens exclusively in low-level jobs, however, and not in high-level jobs. In our model, the introduction of a positive discrimination policy lowers the wages of workers who are supposed to benefit from the policy in both high-level and low-level jobs. This is because a positive discrimination policy makes the positive signal of promotion weaker, while the negative signal of reassignment to the low-level job becomes more significant. In the model by Moro and Norman (2003), all firms obtain the same information about workers' abilities. Therefore, assignment to a specific job does not serve as a signal about worker ability, meaning that the signaling effects that drive results in the current model are not present in their model.

The papers by Athey et al. (2000) and Bjerk (2008) are related to the current paper in that they consider a job hierarchy and study discrimination with respect to promotions. In the model by Athey et al. (2000), workers are employed in a low-level job and need mentoring to enhance their skills and become productive in a high-level job. The skills of a considered worker in the low-level job are enhanced more effectively through mentoring if the share of workers in the high-level job who belong to the same group as the worker is high. As a consequence, if the high-level job is dominated by a group of workers, low-level workers of that group receive more effective mentoring, thereby becoming more productive in the high-level job. Although firms are farsighted and take the effect of promotions on future mentoring into account, it is possible that dominance of a group of workers persists and these workers are more likely to be promoted than the workers of some other group. Bjerk (2008) considers a three-level hierarchy, and new workers begin their career at the lowest level. A worker is promoted if and only if firms have a sufficiently high expectation of the worker's ability. If the average ability in a specific group of workers is low, if the workers belonging to this group cannot send very precise signals of their ability, or if these workers have few chances to signal their ability, even a very able worker from this group needs a long time to upgrade the expectation about his or her ability sufficiently in order to be promoted. Our model differs from the models by Athey et al. (2000) and Bjerk (2008) in that we assume that a worker's employer receives more comprehensive information about the worker's ability than external firms, implying that promotions serve as signals about worker ability. In contrast, Athey et al. (2000) assume that abilities are observable, whereas Bjerk (2008) assumes that there is uncertainty regarding abilities, but learning is symmetric and all firms receive the same information about a worker's ability.

Another related paper is by Milgrom and Oster (1987). They consider a model with two types of workers that they call Visibles and Invisibles. Visibles are workers whose skills are publicly observable. On the contrary, the skill of an Invisible can only be observed by the current employer and becomes commonly known only if the employer decides to promote the Invisible to a high-level job. A consequence of this assumption is that firms have an incentive to "hide" their Invisibles in the low-level job, implying that these workers are discriminated against with respect to promotion decisions. Obviously, one commonality between our model and the model by Milgrom and Oster (1987) is that promotions serve as signals about worker ability. The important difference is that Milgrom and Oster (1987) assume that the signal is perfect, meaning that a worker's ability becomes common knowledge upon promotion. Instead, in our model, the signal is not perfect. Therefore, lowering the promotion standard may actually hurt workers by making the positive signal of promotion weaker.[10]

As we highlighted before, our results crucially depend on the assumption that learning about worker ability is asymmetric in the sense that a worker's employer receives more comprehensive information about the worker's ability than external firms. A few recent papers investigate whether symmetric or asymmetric learning about worker ability is relatively more common in labor markets. Schönberg (2007) develops a formal model that allows for both types of learning. The main finding is that low-ability and high-ability workers are equally likely to switch firms when all firms learn at the same rate about worker ability (i.e., when learning is symmetric), whereas low-ability workers are relatively more likely to move to a different firm when there is asymmetric learning. Schönberg (2007) uses data from the NLSY79 to test the model and she finds evidence consistent with asymmetric learning only for college-educated workers. Pinkston (2009) considers a model with multiple periods. A worker is initially hired by some firm and in each period another firm tries to hire the worker away from the incumbent employer. Each firm receives a private signal about the worker's ability. Importantly, it is assumed that firms engage in a bidding war for the worker's services, which means that firms are able to deduce other firms' signals through the observation of those firms' wage offers. In consequence, when the incumbent employer manages to retain

[10]DeVaro et al. (2012) extend the model by Milgrom and Oster (1987) and they derive some implications that they test with data obtained from the personnel records of a large US firm. In Section 3.3, we explain how our model relates to their empirical findings. Dato et al. (2016) conduct a laboratory experiment to test the models by Waldman (1984) and Milgrom and Oster (1987).

the worker for a long amount of time, the employer becomes ever better informed about the worker's true ability. The worker's wage then approaches the employer's expectation of the worker's productivity which, in turn, becomes an ever more accurate measure of the worker's actual productivity. The latter finding is used to construct an empirical test of the model, which is examined again using data from the NLSY79. The empirical evidence highlights the importance of asymmetric employer learning.

Promotion decisions are often modeled as a tournament in which workers exert a costly effort to perform better than their coworkers and be considered for promotion.[11] A few papers investigate the effects of positive discrimination policies in a tournament setting. In both a theoretical model and an experimental study, Schotter and Weigelt (1992) demonstrate that not only do the workers who are intended to benefit from the policies win the tournament more often (so that their career prospects are improved), but also that efficiency is increased. The intuition for the latter result is that positive discrimination policies tend to make the tournament more equal, thus inducing workers to exert a higher effort.[12] In another set of experiments, Balafoutas and Sutter (2012) show that positive discrimination policies aimed at improving the career prospects of women encourage women to participate in (promotion) tournaments more often instead of working under a piece-rate scheme. Again, it is found that the policy does not entail an efficiency loss. The most important difference between these studies and our paper is that the studies assume wages (i.e., tournament prizes) to be exogenously given, while in our model wages are determined by the firms' competition for workers' services. Wages in our model depend on what external firms learn about a worker's ability from the assignment of the worker to a specific job. As explained previously, wages can therefore be lower when a positive discrimination policy is in place than when there is no such policy, potentially hurting the workers who are intended to benefit from the policy.

[11] The seminal paper on promotion tournaments is by Lazear and Rosen (1981). Konrad (2009) provides an extensive survey of the literature on tournaments and contests.

[12] Similar results are obtained by Fu (2006) in the context of an allpay-auction, which he uses to study race-conscious preferential admissions to college, and by Calsamiglia et al. (2013) in real-effort tournaments between pairs of children. Brown and Chowdhury (2017) study the handicapping of strong participants in professional horseracing, which they find to increase destructive (or sabotage) effort.

3 The basic model

3.1 Description of the model and notation

We consider a model of a competitive labor market with two periods, $\tau = 1, 2$. There are N identical firms and a continuum of workers of (Lebesgue-) measure n; all parties are risk-neutral. There are two different types of job, a low-level job 1 and a high-level job 2. Jobs are indexed by $k = 1, 2$. Each firm has a continuum of jobs of either type that the firm wishes to fill in period τ, and the respective measure for job k is denoted by $M_{k\tau}$.

Workers emanate from two different groups, $l = A, B$, whose measures are n_A and n_B, respectively ($n_A + n_B = n$). If worker j from group l is hired by firm $i \in \{1, ..., N\}$ in period τ and assigned to job k, his or her output is given by

$$y^k_{ijl\tau} = (1 + s_{ij\tau})(c_k + d_k a_{jl}). \tag{1}$$

The worker's ability is denoted by a_{jl} and is initially unknown to all firms and all workers (as, for example, in Holmström 1982). We assume that a_{jl} is continuously and independently distributed. Furthermore, within a group l, the probability density function (pdf) of a_{jl} is the same for all j. It is denoted by f_l and has (full) support $[\underline{a}_l, \overline{a}_l]$, with $\overline{a}_l > \underline{a}_l \geq 0$. The corresponding cumulative distribution function (cdf) is denoted by F_l. $c_k \geq 0$ and $d_k > 0$ are parameters characterizing worker productivity. Following Waldman (1984), we assume that $d_2 > d_1$ (and $c_2 < c_1$), so that output is more responsive to ability in the high-level job. We define $a^e := (c_1 - c_2)/(d_2 - d_1)$ as the ability level at which output is equalized across jobs and we assume $a^e \in (\underline{a}_l, \overline{a}_l)$. Finally, $s_{ij\tau} \in \{0, S\}$ is an indicator variable capturing firm-specific human capital acquired in the first period of employment. Its realization is equal to zero ($s_{ij\tau} = 0$) if the first period is considered or if the second period is considered and worker j has moved to a different firm after the first period. The variable equals $S > 0$ if the second period is considered and the worker continues to work for the same firm as in the first period.

Throughout Section 3, we assume that jobs are "plentiful". In particular, we assume that $M_{k\tau} > n$ for all k and τ, so that all workers could in principle be hired by the same firm and assigned to the same job.[13] In addition, we assume that no positive discrimination policy restricting the firms' decisions is in place. Under this set of assumptions, all the workers are hired and promoted independently of each other. W.l.o.g. we therefore restrict our attention

[13] We introduce slot constraints in Subsection 4.4.

to one representative firm and the firm's workers from one of the groups. We assume that the group has positive measure.

The workers' expected ability, $E[a_{jl}]$, is assumed to be lower than a^e, so that the firm finds it optimal to assign workers to the low-level job 1 in $\tau = 1$.[14] At the end of the first period, the firm observes the output of each worker and then decides which job the workers are assigned to in $\tau = 2$. Other firms (which are also referred to as the "labor market") cannot observe individual outputs, but can observe which job each worker is assigned to at the end of the first period. They use this information to update their ability assessment for the workers. We assume that $\overline{a}_l$ is so high that there is at least one ability level such that the firm wants to promote workers of that ability to the high-level job at the end of the first period.

At the beginning of the second period, other firms attempt to hire workers by making wage offers. It is assumed that all wage offers (including the one from the current employer) are made simultaneously. Each worker is hired by the firm making the highest offer. Ties are broken randomly except for the case in which the current employer is among the firms offering the highest wage. In this case, the worker remains with the current employer. We assume S to be sufficiently high so that, in equilibrium, firms are never successful at hiring workers away from the first-period employer.[15] As in Greenwald (1986) and Waldman (2013), however, there is a (small) probability γ that a worker will switch employers after the first period for exogenous reasons that are unrelated to ability and job assignment; here, the decision to switch employers is taken only after job assignments have been made. As explained in these two papers and in the proof of Lemma 1b), these assumptions eliminate the winner's-curse effect. When a worker switches employers for exogenous reasons, he or she accepts the highest wage offer by the external firms. If the highest offer comes from multiple firms, the worker accepts each offer with the same probability.

Explicit incentive schemes that link pay to performance are not feasible; nor are long-term contracts that bind workers to the firm for both periods. There is no discounting.

The time structure can be summarized as follows: At the beginning of period 1, all firms

[14] The firm's profit would be the same if some group of workers of zero measure were assigned to the high-level job 2 in $\tau = 1$. Here and in some other instances, we disregard deviations for groups of workers of zero measure; see also Assumption 1 in the next subsection.

[15] See Lemma 1 in the subsequent section.

make simultaneous wage offers to all workers and, upon the workers' acceptance decisions, employment relationships begin. Workers are assigned to the low-level job and produce period-1 outputs. At the end of the first period, their current employers observe these outputs and decide about which workers to promote to the high-level job. These decisions are followed by all firms making simultaneous wage offers to all workers for the second period. All newly-hired workers are then assigned to a job, whereas the firms' existing workers remain in their assigned jobs. Finally, the workers produce their period-2 outputs.

3.2 Model solution

As described above, the firm observes the workers' first-period outputs and then decides which of the workers to promote, i.e., the promotion decisions depend on the observed output levels. According to equation (1), there is a unique linear relationship between first-period output and a worker's ability level a_{jl}. This implies that the firm can perfectly infer each worker's ability from the output observation. In the following, we will thus write the promotion decisions as a function of the realized ability levels (instead of the output levels). A promotion rule for some worker is then a mapping from the set of ability levels into the set of jobs.

We simplify the analysis by imposing Assumption 1. The assumption is not very restrictive, as we explain in the paragraph following the assumption's statement.

Assumption 1 *We only consider equilibria with the following properties:*
a) The firm chooses the same promotion rule for all workers from group l.
*b) An external firm offers the same period-*2 *wage to all workers from group* l *who were assigned to job* k *at the end of the first period.*
*c) None of the firms offers a period-*2 *wage to some worker higher than the output that the firm expects the worker to produce if employing the worker.*

Parts a) and b) of Assumption 1 impose symmetry within groups in the sense that all workers from a group are subject to the same promotion rule and offered the same wage by an external firm conditional on their job assignment. These two parts come with very little loss of generality. First, they preclude the possibility that firms treat almost all of the workers with identical observable characteristics in the same way, but display different behavior towards a null set of workers.[16] Since, in the economic literature, it is typically

[16] A null set is a set with measure zero.

assumed that strategies are identical if they entail a different treatment only for a null set of workers, this preclusion is not really restrictive.[17] Second, parts a) and b) of Assumption 1 do not allow firms to divide groups into subgroups of workers, each with positive measure, and to treat workers in the different subgroups differently. Again, this is not really restrictive. The reason is that we could simply extend the number of groups from 2 to any larger number without changing any of our results. Thus, if a firm wishes to divide the workers from a group into two subgroups, and to treat workers from the different subgroups differently, we could simply model this by adding another group with the same ability distribution as the originally considered group. Finally, part c) of the assumption excludes weakly dominated strategies in the sense that firms do not offer wages to workers that are so high that the firms would expect to make a loss when hiring the workers.

Denote by $A_{1l} \subseteq [\underline{a}_l, \overline{a}_l]$ the set of ability levels for which the firm decides to reassign workers from group l to the low-level job in $\tau = 2$, and by $A_{2l} \subseteq [\underline{a}_l, \overline{a}_l]$ the set of ability levels for which workers from group l are promoted to the high-level job. We assume that the firm always promotes a worker to job 2 when it is indifferent between assigning the worker to either job 1 or job 2. Under this assumption, $\{A_{1l}, A_{2l}\}$ is a partition of the set $[\underline{a}_l, \overline{a}_l]$. Furthermore, we denote the external firms' belief regarding $\{A_{1l}, A_{2l}\}$ by $\left\{\tilde{A}_{1l}, \tilde{A}_{2l}\right\}$.

While a worker's first-period employer can infer the worker's ability from the observation of y^1_{ijl1}, the labor market's ability assessment of the worker depends on which job the worker is assigned to at the end of the first period. The same holds for the worker's second-period wage, w_{l2}, which is determined in the following lemma.

Lemma 1 *a) There exists a threshold value $S_1 > 0$ such that, in equilibrium, external firms are never successful at hiring workers away from the first-period employer if $S > S_1$ (unless workers switch firms for exogenous reasons).*

b) If $S > S_1$, then the equilibrium second-period wage for all workers from group l when assigned to job k is given by

$$w_{l2}(k) = \max\left\{c_1 + d_1 E\left[a_{jl} \middle| a_{jl} \in \tilde{A}_{kl}\right], c_2 + d_2 E\left[a_{jl} \middle| a_{jl} \in \tilde{A}_{kl}\right]\right\},$$

where $E[\cdot|\cdot]$ denotes the conditional expectation operator. This wage is offered by the first-period employer and at least one external firm. All other external firms offer a wage smaller than $w_{l2}(k)$.

[17] See, e.g., Ewerhart (2014), footnote 2.

In what follows, we impose the assumption $S > S_1$ and we determine the second-period profit that the first-period employer i earns from employing worker j. When the employer assigns the worker to job k, second-period profit can be stated as

$$\pi_{ijl}(k) := (1-\gamma)\left((1+S)(c_k + d_k a_{jl}) - w_{l2}(k)\right).$$

Obviously, the firm wants to maximize the worker's output (which depends on worker ability), while at the same time wishing to keep wage costs as low as possible. The next proposition characterizes the firm's optimal promotion rule.

Proposition 1 *In equilibrium the following holds:*
a) There exists a threshold value a_l^p such that the firm promotes all workers j from group l at the end of $\tau = 1$ if and only if $a_{jl} \geq a_l^p$, i.e., $A_{1l} = [\underline{a}_l, a_l^p)$ and $A_{2l} = [a_l^p, \overline{a}_l]$.
b) The workers' second-period wages on the two job levels are given by
$w_{l2}(1) = w_{l2}(1, a_l^p) := c_1 + d_1 E[a_{jl} | a_{jl} < a_l^p]$ *and*
$w_{l2}(2) = w_{l2}(2, a_l^p) := c_2 + d_2 E[a_{jl} | a_{jl} \geq a_l^p] > w_{l2}(1, a_l^p)$.
c) The firm's additional second-period profit when promoting worker j (as a function of a_{jl} and a_l^p) corresponds to
$\Delta\pi_l^p(a_{jl}, a_l^p) = (1-\gamma)\left[(1+S)(c_2 - c_1 + (d_2 - d_1) a_{jl}) - (w_{l2}(2, a_l^p) - w_{l2}(1, a_l^p))\right]$. *The threshold value a_l^p is implicitly defined by $\Delta\pi_l^p(a_l^p, a_l^p) = 0$.*
d) Any solution to the condition $\Delta\pi_l^p(a_l^p, a_l^p) = 0$ satisfies $a_l^p > a^e$.

Proposition 1 demonstrates that each worker is promoted at the end of the first period if and only if his or her ability is sufficiently high. As a consequence, promotion serves as a (positive) signal of worker ability, and firms offer higher wages to promoted workers rather than to workers who are reassigned to the low-level job. Because of the wage increase in response to promotion, the firm promotes inefficiently few workers, i.e., the promotion standard a_l^p exceeds the efficient standard of a^e. This replicates the main finding in Waldman (1984). Note that it is possible that the optimal promotion standard a_l^p is not uniquely defined, i.e., the condition $\Delta\pi_l^p(a_l^p, a_l^p) = 0$ may have more than one solution.[18] We return to this issue in the next subsection, when we show how our model can be used to explain worker discrimination.

[18] In the Appendix, we provide a specific example illustrating this possibility. Note that in the original model by Waldman (1984) ability was assumed to be uniformly distributed, in which case the optimal promotion standard is always unique.

To conclude this section, we turn to the beginning of the first period. Since the labor market is competitive, firms are willing to pay a worker a wage such that their total expected profit over both periods is zero. Because of firm-specific human capital and asymmetric learning, the firms that manage to hire a worker in the first period earn a strictly positive profit in the second period. Therefore, they are willing to incur a loss in $\tau = 1$, so that the first-period wage exceeds the expected first-period output.[19] The first-period wage is given by

$$\begin{aligned} w_{l1} &= c_1 + d_1 E\left[a_{jl}\right] \\ &\quad + (1-\gamma) F_l\left(a_l^p\right)\left((1+S)\left(c_1 + d_1 E\left[a_{jl} | a_{jl} < a_l^p\right]\right) - w_{l2}\left(1, a_l^p\right)\right) \\ &\quad + (1-\gamma)\left(1 - F_l\left(a_l^p\right)\right)\left((1+S)\left(c_2 + d_2 E\left[a_{jl} | a_{jl} \geq a_l^p\right]\right) - w_{l2}\left(2, a_l^p\right)\right). \qquad (2) \end{aligned}$$

3.3 Discrimination

Thus far, we have been silent about the issue of discrimination. The model is able to capture both endogenous and exogenous statistical discrimination. Furthermore, by slightly modifying the model, we could also address the situation in which firms discriminate against some workers because of distaste for these workers. We discuss these possibilities in turn.

We begin with *statistical discrimination* that emerges *endogenously*, as in Coate and Loury (1993b). Consider two workers emanating from the different groups A and B and, for simplicity, denote the two workers as workers A and B. Suppose that the two workers are identical ex ante, that is, $f_A = f_B$. Moreover, let there be more than one solution for a_l^p to the condition $\Delta\pi_l^p(a_l^p, a_l^p) = 0$. It is then conceivable that the two workers face different promotion standards a_A^p and a_B^p $(\neq a_A^p)$, although they are identical ex ante. The intuition for this result is the following: If external firms expect a worker's first-period employer to set a rather high promotion standard, they conclude that a promoted worker must be of exceptionally high ability. Therefore, they make very generous wage offers to promoted workers, which in turn makes it optimal for the first-period employer to reassign the worker to the low-level job

[19] One feature of the current model (and the one by Waldman 1984) is that wages can decrease over time, meaning that w_{l1} can be higher than w_{l2}. This result is at odds with real-world compensation schemes which tend to be increasing over time. One way to reconcile the model results with the empirical facts is to assume that workers acquire general human capital (in addition to firm-specific human capital) in the model's first period, which makes them more productive at other firms in the second period, implying that the second-period wage increases.

unless his performance is exceptionally high. In other words, the external firms' expectation of a very high promotion standard may become self-fulfilling in the current model. The same is true for a relatively lower promotion standard, implying that the optimal promotion standard may not be uniquely defined. Assume that worker A faces a higher promotion standard than worker B, $a_A^p > a_B^p$. If both workers are identical both ex ante and ex post, so that they have the same ability $\hat{a}$, and if $\hat{a} \in [a_B^p, a_A^p)$, worker B is promoted to the high-level job, whereas worker A is reassigned to the low-level job. In addition, because the lower promotion standard a_B^p is already inefficiently high, worker A receives a lower total income than worker B. This means that two (ex ante and ex post) identical workers are treated differently and worker A is discriminated against. The latter effect requires that firms correctly anticipate the equilibrium that is played at the end of the first period. Here, it is conceivable that firms use identifiable factors such as the race or sex of a worker to coordinate equilibrium, implying discrimination against workers who are "trapped" in the inefficient equilibrium because of these factors.

When the solution for a_l^p to the condition $\Delta\pi_l^p(a_l^p, a_l^p) = 0$ is unique, the model can still capture *exogenous statistical discrimination.* To see this, consider again two different workers A and B, but assume that the pdfs characterizing the distribution of ability are different, $f_A \neq f_B$. If the ability distributions are different, it is typically the case that workers differ in their value of $c_2 - c_1 + d_2 E[a_{jl}|a_{jl} \geq a_l^p] - d_1 E[a_{jl}|a_{jl} < a_l^p] = w_{l2}(2, a_l^p) - w_{l2}(1, a_l^p)$. In turn, it is optimal for the firm to set different promotion standards for the two workers ($a_A^p \neq a_B^p$), so that there are first-period output levels at which one of the workers is promoted, whereas the other worker is reassigned to the low-level job.

Finally, the model could also be modified to account for the firms' *taste-based discrimination* against workers. The easiest way to incorporate taste-based discrimination into the model is to assume that firm i suffers some disutility $\Delta_{ik} \geq 0$ (per period) when the firm hires a member of some specific group of workers into job k. For the moment, let us continue to assume that all firms are identical so that $\Delta_{ik} := \Delta_k$ for all i. Furthermore, suppose that $\Delta_1 \leq \Delta_2$, meaning that firms care more about what type of workers they hire into the high-level rather than the low-level job. Because the second-period wage upon assignment to job k is determined by the external firms' wage offers and these offers decrease by Δ_k compared to our original model, the second-period wage decreases by Δ_k as well. A direct implication is that the wage spread becomes lower (unless $\Delta_1 = \Delta_2$), whereas the optimal promotion rule stays the same. To understand the latter result, notice that the firm now suffers a relatively greater disutility from promoting the worker of $\Delta_2 - \Delta_1$, while the wage increase upon promotion decreases by $\Delta_2 - \Delta_1$; these effects cancel out so that the promotion decision is the same as in our main model.[20] Finally, with a similar argument as for the second-period wage, the first-period wage is lowered by Δ_1.

The argumentation is a bit more complicated when Δ_{ik} is not constant across firms. In such a situation, it makes sense to assume that, if $\Delta_{i'1} \leq \Delta_{i''1}$ for two different firms $i', i'' \in \{1, ..., N\}$, then we also have $\Delta_{i'2} - \Delta_{i'1} \leq \Delta_{i''2} - \Delta_{i''1}$ (which in turn implies $\Delta_{i'2} \leq \Delta_{i''2}$), with the consequence that firms can be ordered with respect to their discriminatory tastes. To simplify the situation, we assume that $\Delta_{i'k} = \Delta_{i''k}$ for all $i', i'' \in \{2, ..., N\}$ and $k \in \{1, 2\}$, but that $\Delta_{11} < \Delta_{21}$ and $\Delta_{12} - \Delta_{11} < \Delta_{22} - \Delta_{21}$. If there are no slot constraints, as assumed so far, then the least discriminatory firm 1 manages to hire all the workers and pays them a wage depending on the preferences of the other firms. Firm 1 receives a higher payoff from hiring the workers for two reasons. First, there is the obvious effect that the firm suffers lower disutility from hiring the workers. Second, the firm sets a relatively lower promotion standard, i.e., in contrast to the model with identical firms, taste-based discrimination can now affect the promotion rule. Interestingly, taste-based discrimination then leads to relatively more workers being promoted. The reason is simple. When promoting a worker, firm 1 suffers (additional) disutility of $\Delta_{12} - \Delta_{11}$. The wage increase upon promotion, however, is lowered by $\Delta_{22} - \Delta_{21} > \Delta_{12} - \Delta_{11}$. Netting out these effects, firm 1 has a higher incentive to promote the worker, thus leading to a relatively lower promotion standard.[21] Finally, if there exist slot constraints such that the least discriminatory firm is unable to hire all the workers, workers will obviously be hired by multiple firms.

It is to be noted that the three types of discrimination have different implications for workers' career prospects and their wages. Both types of statistical discrimination imply that some workers find it more difficult than others to become promoted, whereas taste-based

[20] Of course, a worker would always prefer not to work in $\tau = 2$ rather than to work and to receive a negative wage. In other words, wages cannot be negative, so that the preceding argument only holds if Δ_k does not exceed the wage in job k that would prevail in the absence of discrimination. Instead, if Δ_k is higher than this wage, taste-based discrimination affects the promotion rule.

[21] Setting a lower promotion standard could be detrimental if the standard becomes inefficiently low. An inefficiently low standard could always be ruled out by assuming that the firms' preferences do not differ too strongly.

discrimination does not always have an effect on the promotion rule or may even lead to more of the disliked workers being promoted. To understand the effect of discrimination on wages, consider a worker who faces a relatively higher promotion standard because of (some type of statistical) discrimination, but still manages to become promoted. In the case of endogenous statistical discrimination, it is assumed that the ability distribution is the same for all workers. Thus, by being promoted, the worker signals a high ability and is rewarded with a relatively high wage offer. In contrast, in the case of exogenous statistical discrimination, the ability distributions for different types of workers are assumed to be different. As a result, a worker can be offered a lower wage even if he or she faces a higher promotion standard. Similarly, our simple model of taste-based discrimination predicts lower wages for workers who are discriminated against. In light of the frequent observation that people who are perceived to be discriminated against also receive lower wages upon promotion,[22] these results allow the conclusion that exogenous statistical discrimination and taste-based discrimination are potentially more important than endogenous statistical discrimination.

Recent findings by DeVaro et al. (2012) underscore the argument that exogenous statistical discrimination could be responsible for the relatively bad career prospects of some groups of people. They use data obtained from the personnel records of a large US firm and they find that white workers face lower promotion standards than non-white (i.e., black, Hispanic, or Asian) workers. When tasks are more variable across hierarchical levels, it is observed that the disadvantage of non-white workers relative to white workers is mitigated. Bjerk (2008) argues that non-white workers tend to attend relatively worse schools than white ones, implying that their average ability when entering the labor market is lower. In addition, the quality of schools that non-whites attend is probably more variable than the quality of schools that whites attend, resulting in a higher variance of the non-whites' abilities.[23] To account for both these effects, we assume that the ability of white workers is uniformly

[22] For instance, the study by Bertrand et al. (2014) finds that women appointed to the boards of publicly traded Norwegian companies earned about 38% less than their male counterparts (before the passing of the law mandating a 40-percent representation of each gender).

[23] A similar argument is advanced by Aigner and Cain (1977). They argue that certain labor market signals are less reliable indicators of performance for black workers rather than white workers, with the consequence that the employers' ability estimates for black workers are less precise. They use this argument to explain why black workers are discriminated against. While their explanation relies on the firms' risk aversion, our results hold under risk neutrality.

distributed on $[\underline{a}_w, \overline{a}]$, whereas the ability of non-white workers is uniformly distributed on $[\underline{a}_{nw}, \overline{a}]$, with $\underline{a}_{nw} < \underline{a}_w$. Given the uniform distribution, the optimal promotion standards can be determined as

$$a_w^p = \frac{d_2\overline{a} - d_1\underline{a}_w + 2S\,(c_1 - c_2)}{(d_2 - d_1)\,(1 + 2S)}$$

and

$$a_{nw}^p = \frac{d_2\overline{a} - d_1\underline{a}_{nw} + 2S\,(c_1 - c_2)}{(d_2 - d_1)\,(1 + 2S)},$$

implying a relatively higher promotion standard for the non-white workers ($a_w^p < a_{nw}^p$). Notice that

$$a_{nw}^p - a_w^p = \frac{d_1\,(\underline{a}_w - \underline{a}_{nw})}{(d_2 - d_1)\,(1 + 2S)}.$$

When tasks become more variable across hierarchical levels, it is likely that d_2 increases and/or d_1 decreases (of course, c_1 and c_2 can change as well). In either case, $a_{nw}^p - a_w^p$ becomes lower, mitigating the disadvantage of the non-white workers relative to the white ones, as found by DeVaro et al. (2012).

4 Positive discrimination policy

We now introduce a positive discrimination policy into the model that is aimed at improving the career prospects of people who are discriminated against. The policy that we consider sets a quota $\alpha \in [0, 1]$ for the workers of group A in the high-level job, meaning that workers belonging to group A must fill a share of at least α of the positions in the high-level job. If the promotion standard is not unique, even a small policy intervention may induce firms and workers to switch from one equilibrium to another, implying a substantial change in the promotion standard. We avoid such difficulties by focusing on the effects of the policies conditional on a specific equilibrium being played. The easiest way to justify this procedure is to come up with conditions that guarantee that the solution for a_l^p to the condition $\Delta\pi_l^p\,(a_l^p, a_l^p) = 0$ is unique. This is assumed in the following.[24] We write a_{li}^p if we refer to the promotion standard of a specific firm i.

[24] As mentioned before, if abilities are uniformly distributed, as assumed, e.g., by Waldman (1984), then the solution for a_l^p to the condition $\Delta\pi_l^p\,(a_l^p, a_l^p) = 0$ is always unique.

4.1 Workers who are already employed when the policy is introduced

Our analysis is divided into two parts. We begin by studying the effects of the policy on workers who have already been hired and who are in the first half of their career (i.e., the first period) when the policy is introduced and which we term the policy's short-run effects. Thereafter, we investigate the policy's long-run effects on workers who begin their career after the policy was introduced. Studying the short-run effects of the policy makes sense for a number of reasons. First, it may take a considerable amount of time until the ultimate effects of the policy are borne out. Firms may react to the introduction of the policy by changing the composition of their workforce, potentially hiring new workers from specific worker groups, while separating from other workers. These types of decisions are lengthy, meaning that the transition from the old to the new equilibrium may take quite long so that it is important to study the effects of the policy during the transitional period. Second and related to the first point, policymakers often face a lot of pressure to reverse policies, in particular when these policies fail to achieve their goals in the short run. To understand whether policymakers should succumb to the pressure or whether they should stick with their policies, it is important to know whether and how the policy's short-run effects differ from the effects in the long run. Third, several countries have lately introduced policies aimed at facilitating career advancement for certain groups of people and researchers have started to investigate the effects of these policies. Given the recent introduction of the policies, it is likely that the considered countries are situated in the transitional period in which firms are still adapting to the policies' introduction. Formally studying the short-run effects of the policy is therefore important to derive some implications that can be contrasted with the existing empirical evidence. Fourth, considering the short-run effects is also beneficial for pedagogical reasons. Understanding the short-run effects of the policy makes it easier to understand the policy's long-run effects, since the short-run effects operate in the long run as well.

In the short run, two variables are not affected by the policy. First, the workers' period-1 wages are fixed because these wages have already been determined before the policy was put in place. Second, each firm's composition of the workforce is fixed as well because period-1 hiring decisions have already been made. Regarding the latter, we make the assumption that firm i originally managed to hire a continuum of workers emanating from group l of measure $n_{li} > 0$. We consider a symmetric situation by assuming that $n_{li} = \frac{n_l}{N}$ for all $i \in \{1, ..., N\}$ and $l \in \{A, B\}$. The interpretation is that all N firms made the same wage offers to all workers from the same group l before the quota was introduced, with the consequence that each firm managed to hire the same "number" of workers.[25] This symmetry assumption implies that firms manage to hire the very same measures of workers that they lose to external firms at the beginning of period 2, so that the newly-hired workers can simply replace the leaving workers. The following proposition characterizes the short-run effects of the policy.

Proposition 2 *Suppose that all firms hired a continuum of workers from groups A and B of measures $\frac{n_A}{N}$ and $\frac{n_B}{N}$, respectively, before the positive discrimination policy was introduced. Further, suppose that $S > S_2$, where $S_2 > 0$ is a threshold value. Then there is α_1 such that, if $\alpha < \alpha_1$, an equilibrium with the following properties exists after the introduction of the policy:*

a) Each firm promotes all workers from group $l \in \{A, B\}$ with ability $a_{jl} \geq a_l^q$ to the high-level job at the end of $\tau = 1$ and reassigns all other workers to the low-level job, where $a_A^q \, (\leq a_A^p)$ and $a_B^q \, (\geq a_B^p)$ are the new promotion thresholds after the quota α was introduced.

b) At the beginning of $\tau = 2$, all firms offer the same wage $w_{l2}(k, a_l^q)$ to each worker from group l who was assigned to job k. The wages for the two jobs are given by

$$\begin{aligned} w_{l2}(1, a_l^q) &= c_1 + d_1 E\left[a_{jl} | a_{jl} < a_l^q\right] \text{ and} \\ w_{l2}(2, a_l^q) &= c_2 + d_2 E\left[a_{jl} | a_{jl} \geq a_l^q\right] > w_{l2}(1, a_l^q). \end{aligned}$$

c) With probability $1 - \gamma$, a worker accepts the offer of his first-period employer. With probability γ, the worker switches the employer for exogenous reasons. In this case, the worker randomly selects each of the external firms' offers with the same probability $1/(N-1)$.

d) The period-1 employers always manage to retain their workers unless workers switch firms for exogenous reasons. Each firm loses a continuum of workers from group l who were assigned to job 1 of measure $\gamma F_l(a_l^q) \frac{n_l}{N}$. At the same time, each firm manages to hire a continuum of workers from group l who were assigned to job 1 at other firms of the same measure $\gamma F_l(a_l^q) \frac{n_l}{N}$, and replaces the departing workers from group l in job 1 with the newly-hired ones from the same group and job. The same holds for job 2, where the corresponding measure is $\gamma(1 - F_l(a_l^q)) \frac{n_l}{N}$.

[25] Notice that such a symmetric equilibrium always exists in the model from Section 3.

The most important implication of Proposition 2 is that the introduction of the policy leads to changes in the two groups' promotion standards; for group A, the standard is lowered from a_A^p to a_A^q; for group B, the standard is increased from a_B^p to a_B^q. Intuitively, because of the quota it is possible that the first-period employer is ordered to promote more workers from group A or fewer workers from group B than he or she voluntarily would. It is then optimal for the employer to promote the next-best workers from group A or to demote the worst workers from group B, implying the described changes in the promotion standards.[26]

Except for the change in the promotion standards, the results are similar to those of the basic model. Period-2 wages are again determined by competition among the firms for the workers' services. Because workers are more valuable to the current employer than to external firms, the employer manages to retain workers unless they decide to switch firms for exogenous reasons. And as explained before, since firms are initially of equal size, firms hire the very same measures of workers that they lose to external firms, meaning that the newly-hired workers can simply replace the leaving workers. The condition $\alpha < \alpha_1$ is important for this latter finding and also for the period-2 wage offers. The condition implies $E\left[a_{jA}|a_{jA} \geq a_A^q\right] \geq a^e$, which means that the promotion standard for group A never becomes so low in response to the policy introduction that firms expect promoted workers from this group to be more productive in the low-level rather than the high-level job.

In the remainder of this subsection, we will take a closer look at the effects of the policy on workers' wages. We concentrate on the workers from group A, who are supposed to benefit from the policy.

Proposition 3 *Suppose that the positive discrimination policy affects short-run behavior as described in Proposition 2.*

a) Then $w_{l2}(k, a_l^q)$ is strictly increasing in a_l^q.

b) When the policy lowers the promotion standard for group A from a_A^p to a_A^q, workers with

[26] Estevan et al. (2014) make similar arguments in the context of college admission. They study the Texas top ten percent policy, which guarantees Texas students who graduated in the top ten percent of their high school class automatic admission to all state-funded colleges. The rule makes it easier for students from high schools with many disadvantaged students to be admitted to college, which is equivalent to a decrease in the educational standard that these students must meet. Interestingly, Estevan et al. (2014) find that the rule induced some students to switch strategically to a different high school to improve their chances of being admitted to college.

ability $a_{jA} \in [\underline{a}_A, a_A^q) \cup [a_A^p, \bar{a}_A]$ *receive a lower second-period wage and workers with ability* $a_{jA} \in [a_A^q, a_A^p)$ *receive a higher wage.*

The positive discrimination policy has two effects on the second-period payoff of a worker from group A who is favored by the policy. First, the worker is more likely to be promoted and obtain a wage increase, which obviously benefits the worker. Second, the positive signal of promotion to the high-level job becomes weaker, whereas the negative signal of being reassigned to the low-level job becomes stronger. If the worker is promoted, external firms may question the worker's ability and may believe that the worker was promoted only because of the positive discrimination policy. If the worker is not promoted, external firms may believe that the worker's ability must be extremely low because he or she was not promoted in spite of the positive discrimination program. Because external firms interpret the job-assignment signal differently when a positive discrimination policy is in place, their wage offers also differ. In particular, both the wage for a promoted worker and a worker reassigned to the low-level job decrease, as shown in part a) of Proposition 3. This obviously hurts workers. The second part of the proposition demonstrates that the introduction of the positive discrimination policy leaves those workers who either have a very high ability so that they would have been promoted even without the policy, or who have such low ability that they are not considered for promotion even when the policy is in place, worse off. In contrast, workers in the middle of the ability distribution benefit from the policy because those workers are promoted if and only if firms are bound to the positive discrimination policy.

It is possible that the negative effects on workers of either low or high ability outweigh the positive effects on the workers of middle ability, so that a worker's expected payoff may actually decrease. As will be shown in Proposition 4, this happens only if the promotion standard is sufficiently below the efficient level that the inefficiency due to a standard that is too low is larger than the promotion signaling distortion which causes an inefficiently high standard.

Proposition 4 *Suppose that the positive discrimination policy affects short-run behavior as described in Proposition 2. Denote by* $W_{A2}(a_A^q)$ *the expected second-period wage of a worker from group A as a function of* a_A^q *and define* $\hat{a} \in [\underline{a}_A, a^e)$ *implicitly by* $E[a_{jA} | a_{jA} \geq \hat{a}] = a^e$.

a) W_{A2} *has a global maximum at* a^e.

b) There exists a threshold $\hat{a}^q \in (\hat{a}, a^e)$ *such that* $W_{A2}(a_A^q) < W_{A2}(a_A^p)$ *for all* $a_A^q \in (\hat{a}, \hat{a}^q)$.

4.2 Workers who begin their career after the policy is introduced

In this subsection, we investigate the long-run effects of the policy. Compared to the previous subsection, the effects of the policy may be altered for two reasons. First, the period-1 wage is no longer fixed and will be changed due to the policy. Second, firms may also have an incentive to change the composition of their workforce. For instance, by hiring more workers from group A into the low-level job at the beginning of $\tau = 1$, firms find it easier to satisfy the quota in the high-level job in $\tau = 2$, thereby affecting the promotion standards for the two groups of workers.

In order to determine an equilibrium of the game, we start the analysis by considering an auxiliary program. For ease of notation, define

$$Y_{l1} := c_1 + d_1 E\left[a_{jl}\right] \text{ and}$$
$$Y_{l2}\left(a_l^p\right) := F_l\left(a_l^p\right)\left(c_1 + d_1 E\left[a_{jl} | a_{jl} < a_l^p\right]\right) + \left(1 - F_l\left(a_l^p\right)\right)\left(c_2 + d_2 E\left[a_{jl} | a_{jl} \geq a_l^p\right]\right)$$

as the expected outputs of a worker from group l in the two periods when the worker's employer sets the promotion standard a_l^p (excluding firm-specific human capital). We also set $q_{Ai} := n_{Ai}/n_{Bi}$ and $q_{Bi} := n_{Bi}/n_{Ai}$ (and thus $q_{Ai} = 1/q_{Bi}$), where n_{li} again represents the measure of workers from group l hired by firm i. Furthermore, as in the proof of Proposition 2, we define $q_{Ai}^p(a_A^p, a_B^p) := \frac{n_{Ai}\cdot(1-F_A(a_A^p))}{n_{Ai}\cdot(1-F_A(a_A^p))+n_{Bi}\cdot(1-F_B(a_B^p))}$ as the share of promoted workers from group A in firm i. The auxiliary program (AP_i) of firm i corresponds to the maximization of the expected average output per worker who begins to work for the firm under the constraint implied by the quota α and the equilibrium constraints that determine the promotion thresholds.[27]

$$\underset{q_{Bi}, a_{Ai}^p, a_{Bi}^p}{Max} \sum_{l \in \{A,B\}} \frac{1}{1+q_{li}} \left[Y_{l1} + \left(1 + (1-\gamma) S\right) Y_{l2}\left(a_{li}^p\right)\right]$$

s.t.

$$q_{Ai}^p(a_{Ai}^p, a_{Bi}^p) > \alpha, \quad \Delta\pi_A^p(a_{Ai}^p, a_{Ai}^p) = \Delta\pi_B^p(a_{Bi}^p, a_{Bi}^p) = 0,$$

or

$$q_{Ai}^p(a_{Ai}^p, a_{Bi}^p) = \alpha, \quad \Delta\pi_A^p(a_{Ai}^p, a_{Ai}^p) \leq 0 \leq \Delta\pi_B^p(a_{Bi}^p, a_{Bi}^p),$$
$$\alpha \cdot \Delta\pi_A^p(a_{Ai}^p, a_{Ai}^p) + (1-\alpha) \cdot \Delta\pi_B^p(a_{Bi}^p, a_{Bi}^p) = 0.$$

[27] Note in the following optimization problem that $n_{li}/(n_{Ai} + n_{Bi}) = 1/(1 + q_{li})$ for $l \in \{A, B\}$.

In the first period, a worker from group l is expected to produce output equal to Y_{l1}. In the second period, the worker produces an expected output of $(1+S)\,Y_{l2}\,(a_{li}^{p})$ when staying with the firm (which happens with probability $(1-\gamma)$) and of $Y_{l2}\,(a_{li}^{p})$ when moving to one of the other firms (which happens with probability γ). The constraints ensure that firm i does not wish to deviate from the optimal promotion standards in the presence of the positive discrimination policy. We explain this in detail in the proof of part c) of Proposition 5.[28]

In the following, we denote the solution to the program by $q_B^*\,(\alpha)$, $a_A^{p*}\,(\alpha)$ and $a_B^{p*}\,(\alpha)$. On the basis of this solution, we are able to characterize an equilibrium when the positive discrimination policy is present. To simplify the exposition, we assume that $n_B/n_A = q_B^*\,(\alpha)$. In words, the ratio of workers from the two groups in the total population is equal to the corresponding ratio solving (AP_i). As we show, in equilibrium all firms want to hire workers from the two groups such that their ratio is $q_B^*\,(\alpha)$. If n_B/n_A were not equal to $q_B^*\,(\alpha)$, some firms would not be able to do so; instead, these firms would be segregated in that they would manage to hire only workers from one of the groups. This would make the analysis more complicated without adding any important new insights. The assumption $n_B/n_A = q_B^*\,(\alpha)$ could be justified by arguing that the workers who are "left over" when all firms hire workers from the two groups, such that their ratio is $q_B^*\,(\alpha)$, leave the industry to work in another industry. We further simplify the exposition by assuming $E\,[a_{jA}|a_{jA} \geq a_A^{p*}] \geq a^e$, implying that external firms expect workers from group A who are promoted by their original employer to be relatively more productive in the high-level than in the low-level job. This is similar to what we assumed in Proposition 2. A sufficient condition for $E\,[a_{jA}|a_{jA} \geq a_A^{p*}] \geq a^e$ to hold is that α is not too high and the workers from group B are not much more productive than the workers from group A, i.e., $Y_{B1} - Y_{A1}$ and $Y_{B2}\,(a) - Y_{A2}\,(a)$ should not be too high. If productivity differences are sufficiently low, the measure of hired workers emanating from group A is sufficiently large that workers from group A are promoted only if their ability is relatively high.

Proposition 5 *Denote by q_B^*, $(q_A^* = 1/q_B^*,)$ a_A^{p*} and a_B^{p*} the solution to problem (AP_i). Assume that $n_B/n_A = q_B^*$ and $E\,[a_{jA}|a_{jA} \geq a_A^{p*}] \geq a^e$. If $S > S_2$, an equilibrium with the following properties exists.*

[28] Notice that either the condition $\Delta\pi_A^p(a_{Ai}^p, a_{Ai}^p) \leq 0$ or $\Delta\pi_B^p(a_{Bi}^p, a_{Bi}^p) \geq 0$ could be removed from the program (AP_i) because it is implied by the remaining two conditions of (AP_i).

a) At the beginning of $\tau = 1$, $\hat{N} \in \{3, ..., N\}$ of the firms offer the same wage w_{l1} to every worker from group $l \in \{A, B\}$. Using $\ell \in \{A, B\} \setminus \{l\}$ this wage is given by

$$w_{l1}^{*} = \frac{1}{1+q_l^*}\left(Y_{\ell 1} + q_l^* Y_{l1} + \left(1 + (1-\gamma) S\right) Y_{\ell 2}\left(a_\ell^{p*}\right) + \left(q_l^*\left(1-\gamma\right) S - 1\right) Y_{l2}\left(a_l^{p*}\right)\right).$$

The remaining $N - \hat{N}$ firms do not make any wage offers.

b) Each worker accepts each of the highest offers with the same probability $\frac{1}{\hat{N}}$. Thus, in $\tau = 1$, each of the $\hat{N}$ firms hires a continuum of workers of measure $\hat{n}_l := n_l/\hat{N}$ from group $l \in \{A, B\}$ and assigns all these workers to the low-level job.

c) At the end of $\tau = 1$, each firm promotes all workers from group l with ability $a_{jl} \geq a_l^{p*}$ to the high-level job and reassigns all other workers to the low-level job.

d) At the beginning of $\tau = 2$, all $\hat{N}$ firms offer the same wage $w_{l2}\left(k, a_l^{p*}\right)$ to each worker from group l who was assigned to job k. The wages for the two jobs are given by

$$\begin{aligned} w_{l2}\left(1, a_l^{p*}\right) &= c_1 + d_1 E\left[a_{jl} | a_{jl} < a_l^{p*}\right] \text{ and} \\ w_{l2}\left(2, a_l^{p*}\right) &= c_2 + d_2 E\left[a_{jl} | a_{jl} \geq a_l^{p*}\right] > w_{l2}\left(1, a_l^{p*}\right). \end{aligned}$$

The remaining $N - \hat{N}$ firms do not make any wage offers.

e) With probability $1 - \gamma$, a worker accepts the offer of his first-period employer. With probability γ, the worker switches the employer for exogenous reasons. In this case, the worker randomly selects each of the highest offers by the external firms with the same probability $1/(\hat{N} - 1)$.

f) The period-1 employers always manage to retain their workers unless workers switch firms for exogenous reasons. Each firm loses a continuum of workers from group l who were assigned to job 1 of measure $\gamma F_l\left(a_l^{p*}\right)\hat{n}_l$. At the same time, each firm manages to hire a continuum of workers from group l who were assigned to job 1 at other firms of the same measure $\gamma F_l\left(a_l^{p*}\right)\hat{n}_l$, and replaces the departing workers from group l in job 1 with the newly-hired ones from the same group and job. The same holds for job 2, where the corresponding measure is $\gamma\left(1 - F_l\left(a_l^{p*}\right)\right)\hat{n}_l$.

The intuition for the proposition is the following. By changing the measures of hired workers from the two groups, firms can affect the promotion standards. For instance, if the measure of hired workers from group A is high, while the measure of hired workers from group B is low, a firm finds it relatively easy to satisfy the quota constraint, so that the promotion thresholds do not deviate strongly (or at all) from the standards derived in the model without the positive discrimination policy. In contrast, if the measure of hired workers from group A is low, while the measure of hired workers from group B is high, the quota constraint becomes much more difficult to satisfy, implying bigger changes in the promotion thresholds. In the equilibrium, firms make hiring decisions such that the ratio of the measures of hired workers from the two groups leads to optimal promotion standards in the sense that the objective function from (AP_i) is maximized.

Another result is that a worker's total expected wage no longer necessarily equals the worker's expected output. In the basic model from Section 3, all hiring decisions were made independently of each other, so that a firm earned zero expected profit from hiring any worker. Because of the positive discrimination policy, hiring decisions are not independent anymore, implying that the zero-profit condition holds on the aggregate level (i.e., for the entirety of workers a firm employs), but not necessarily on the individual level. In the equilibrium characterized in Proposition 5, workers from the two different groups receive exactly the same expected total wage, meaning that all gains from trade are distributed evenly among the workers from the different groups. It is not clear whether gains from trade could be distributed in a different way, i.e., whether an equilibrium with the same properties as that from Proposition 5, but with a different choice of period-1 wages, exists. The period-1 wages described in the proposition lead to a stable solution in the sense that firms make sure that none of the other firms has an incentive to outbid the firm at the beginning of the first period. While a zero-profit constraint on the firm-level could also be fulfilled for other combinations of period-1 wages for the two groups, other wages may tempt firms to deviate from the described equilibrium.

A simple numerical example illustrates this point. Suppose that $q^* = 1$ and a firm wishes to hire a continuum of workers from the two groups of measure 5, respectively. Let the expected output per unit of measure be 6 for group A and 4 for group B, so that the total output of all the workers equals 50. Furthermore, suppose that another firm wants to hire a continuum of workers from group A of measure 1 and a continuum of workers from group B of measure 5 and that the expected output per unit of measure is 4.8 for both groups in this case, resulting in total output of 28.8. Now if the first firm does not distribute the gains from trade evenly among the workers from the two groups, it risks losing some of the workers. To illustrate this, suppose that all workers receive a total expected wage equal to their expected output, that is, all workers from group A receive a total expected wage of 6 per unit of measure, whereas those from group B receive just 4. Then the second firm could overbid the first firm by offering slightly above 6 per unit of measure for the workers from group A and slightly above 4 for the workers from group B. The firm's profit would be close to $28.8 - 26 = 2.8$, so that the second firm would gain from hiring the workers away from the first firm. Instead, if the first firm distributed the gains from trade evenly among the two groups of workers and paid both groups a total expected wage of 5 per unit of measure, the second firm would never have an incentive to hire workers away from the first one.

In the next step, we take a closer look at the effects of the positive discrimination policy on workers' wages. We find that the long-run effects of the policy differ substantially from the corresponding short-run effects. In particular, under some additional conditions, the targeted workers always benefit from the policy regardless of their ability level.

Proposition 6 *Suppose that the equilibrium from Proposition 5 is played. Furthermore, assume that any solution to (AP_i) has $q_A^* > 0$ and $q_B^* > 0$ and that $Y_{A1} \geq Y_{B1}$. Then there is another threshold-value S_3 such that, if $S > S_3$, any worker j of group A receives a higher total wage than in the situation without the policy (where total wage refers to the sum of first and second-period wages when conditioning on the worker's ability level).*

With the same reasoning as in Proposition 3, for a worker of group A the second-period wage, conditional on the job level, always decreases when the promotion standard is lowered due to the positive discrimination policy. In contrast, the first-period wage may well increase. Three effects are at work. First, the change in the promotion standard changes the worker's expected second-period compensation, which in turn affects the firms' willingness to pay for the worker in the first period. Second, as explained before, firms make hiring decisions such that the ratio of the measures of hired workers from the two groups leads to optimal promotion standards in the sense that the objective function from (AP_i) is maximized. This increases the workers' expected outputs and the firms' willingness to pay for workers' services. Third, as also explained, hiring decisions are now interdependent with the consequence that all workers' expected total wages are equalized.

As Proposition 6 indicates, the increase in the period-1 wage can be sufficiently strong such that the effects of the policy on workers' period-2 wages are outweighed and workers receive higher total wages regardless of their ability. The conditions in the proposition ensure that the period-1 wage for a group-A worker responds more strongly to changes in S when the positive discrimination policy is in place rather than when there is no such policy. As a consequence, when S is sufficiently large, all workers from group A benefit from the introduction of the policy.

When all workers from group A benefit from the policy, it is obvious that they also benefit from the policy in expectation, that is, before their ability is known. The corresponding result is described in Proposition 7, which shows that the requirements are somewhat weaker than those in Proposition 6. The effects that drive the results are the same as highlighted in the discussion of Proposition 6.

Proposition 7 *Suppose that the equilibrium from Proposition 5 is played. Furthermore, assume that any solution to (AP_i) has $q_A^* > 0$ and $q_B^* > 0$. Then any worker j of group A receives a higher expected total wage than in the situation without the policy (where expected total wage refers to the sum of first-period and expected second-period wages).*

4.3 Effect on people disadvantaged by the policy

In this subsection, we briefly address the effects of the positive discrimination policy on workers who are disadvantaged by the policy and who therefore face a relatively higher promotion standard, as explained before. Using our previous results, it is clear that some of the workers who are supposed to be disadvantaged by the policy may actually be better off, at least in the short run. As shown in Proposition 3, extremely able workers who manage to become promoted in spite of the policy receive a higher second-period wage because the positive signal of being promoted grows stronger. The same holds for rather unable workers who are not promoted even if there is no positive discrimination policy that makes it more difficult for them to move up the corporate ladder. In contrast, workers who are promoted if and only if no positive discrimination policy is in place are typically worse off. Given that the first-period employer promotes inefficiently few workers, a further increase in the promotion standard reduces expected output. In expectation, workers therefore suffer and receive a lower total wage. This may explain why, in practice, members of an initially advantaged group rarely lobby to have positive discrimination policies enacted, although some of them may actually be better off when such policies are introduced.

4.4 Slot constraints in the high-level job

Up to this point, we imposed the assumption $M_{k\tau} > n$ for all k and τ, so that all workers could in principle be hired by the same firm and assigned to the same job. An implication of this assumption is that each of the firms is able to promote any number of its first-period employees. In practice, however, positions on higher levels of the firms' hierarchies are often scarce. In this subsection, we therefore discuss the impact on the model of slot constraints in the high-level job, that is, we now disregard the assumption $M_{2\tau} > n$. Interestingly, introducing slot constraints into our model does not change any of our results. In the preceding analysis, we started by addressing the basic model without a positive discrimination policy and we proceeded by investigating the policy's short-run and long-run effects. In the remainder of this subsection, we study these three situations one after the other and we explain why our previous results are robust to considering slot constraints in the high-level job.

One finding of the basic model was that firms choose inefficiently high promotion standards for both groups A and B. Consider some firm i and assume that the firm hired a continuum of workers emanating from groups A and B of measures $n_{Ai} > 0$ and $n_{Bi} > 0$, respectively, at the beginning of $\tau = 1$. Denote by a^p_{li} the firm's optimal promotion standard for group l, as determined in Proposition 1. Then, if $n_{Ai}\left(1 - F_A\left(a^p_{Ai}\right)\right) + n_{Bi}\left(1 - F_B\left(a^p_{Bi}\right)\right) \leq M_{22}$, the slot constraint would not really restrict the firm's decision since the firm would not want to promote more workers than allowed by the constraint. On the contrary, if $n_{Ai}\left(1 - F_A\left(a^p_{Ai}\right)\right) + n_{Bi}\left(1 - F_B\left(a^p_{Bi}\right)\right) > M_{22}$, the slot constraint would be violated and the firm forced to promote fewer workers than it would normally want to do. Promoting fewer workers is equivalent to raising (at least one of) the promotion standards, and, since the standards are already inefficiently high, would lead to a reduction in the attainable surplus. Notice that firm i can affect n_{li} by deciding about how many workers to make an offer to at the beginning of $\tau = 1$. Hence, firms could structure their period-1 wage offers such that n_{Ai} and n_{Bi} are sufficiently small that $n_{Ai}\left(1 - F_A\left(a^p_{Ai}\right)\right) + n_{Bi}\left(1 - F_B\left(a^p_{Bi}\right)\right) \leq M_{22}$ holds. In other words, firms could deliberately decide to stay small in order not to be affected by the slot constraint. Given that there is competition among firms and an "effective" slot constraint leads to a reduction in surplus, only the small firms are able to compete effectively for workers' services, implying that, in equilibrium, all the active firms decide to stay small and the slot constraint has no effect on promotion decisions.

Consider now the short-run effects of the positive discrimination policy. Our main argument in Subsection 4.1 was that the policy changes the promotion standards; for group A the standard is lowered, whereas for group B the standard is increased. The slot constraint may in principle affect the size of these changes. To understand this, suppose that the condition $n_{Ai}(1 - F_A(a^p_{Ai})) + n_{Bi}(1 - F_B(a^p_{Bi})) = M_{22}$ holds before the policy is introduced, that is, firm i made hiring decisions such that the slot constraint is just binding at the optimal promotion standards. If the firm wishes to react to the policy by lowering the promotion standard for group A and, thus, promoting more workers from group A, the slot constraint could be violated. This means that a decrease in the promotion standard for group A potentially requires a sufficiently large increase in the promotion standard for group B to meet the slot constraint. The slot constraint may therefore affect the exact values of the promotion standards. The qualitative predictions of the analysis from Subsection 4.1, however, are not affected by the slot constraint. Each firm still tends to react to the policy by lowering the promotion standard for group A and increasing the standard for group B.

Finally, the slot constraint does not affect our conclusions regarding the policy's long-run effects. The argument is essentially the same as for the basic model. In Proposition 5, it was shown that there is an optimal ratio of group sizes and that firms hire workers from the two groups in accordance with that ratio at the beginning of $\tau = 1$. The absolute sizes of the two groups of workers that firms employ, however, were indeterminate and depended on the number of active firms in the market. So if there is a slot constraint in the high-level job, which, if binding, would entail less efficient decisions on the part of the firms, an equilibrium is played in which firms are sufficiently small (or sufficiently many firms are active) such that the slot constraint simply does not restrict the firms' decisions. Put differently, the equilibrium is the same as in Proposition 5, just with a sufficiently large $\hat{N}$.

5 Empirical implications of the model

Several empirical implications can be derived from the model. We studied the short-run and long-run effects of the positive discrimination policy and we begin by discussing the implications that are common to both these studies. In general, the policy could have an effect on the firms' hiring decisions, their promotion decisions, and their wage offers to the workers. When we studied the policy's short-run effects, the initial hiring decisions had already been made, so we start by discussing the effects of the policy on promotions and period-2 wage offers. We demonstrated that the introduction of the policy leads to changes in the promotion standards for the different groups of workers (Proposition 2). A direct implication is that the policy's introduction changes the employment of people from an initially disadvantaged group and an initially advantaged group in high-level jobs. While more people from the former group are promoted to the high-level job, there are fewer people from the latter group in the very same job. In line with this prediction, Holzer and Neumark (2000), surveying the literature on positive discrimination programs, conclude that these "programs redistribute employment [...] from white males to minorities and women" (p. 558). Recent studies such as Kurtulus (2012) and Bertrand et al. (2014) underscore this observation. Kurtulus (2012) finds that the share of minorities and women who are employed in high-paying skilled jobs in the US grew more between 1973 and 2003 in firms that were subject to affirmative action regulations than in firms that were not. Bertrand et al. (2014) study the effects of a law that was passed in Norway in 2003, which mandates a forty-percent representation of each gender on the boards of publicly traded companies. They observe that many firms changed their status to private after 2003 to be exempt from the law. The remaining firms significantly increased the number of female directors on the board, but only after the introduction of severe sanctions for noncompliance.

An important implication of the model is that the introduction of the policy lowers the wages that workers who are advantaged by the policy earn later in their careers, and increases the wages for workers who are disadvantaged by the policy when controlling for job level (Proposition 3). The starting wage of workers who began their working careers before the policy was implemented is unaffected by the policy. This, in turn, implies that the wage increase upon promotion for these workers is changed by the positive discrimination policy. For workers who are advantaged by the policy, the post-promotion wage decreases, while the starting wage is fixed so that the wage increase upon promotion also decreases. For workers who are disadvantaged by the policy, the effect is the opposite and the wage increase upon promotion increases. Note that this is exactly what Bertrand et al. (2014) find in their analysis of Norwegian companies. They show that election to the board of a company entails a substantial financial reward for the elected worker. After the board quota was introduced, this reward fell for women (from 9.4% percent of annual earnings to 8%), while it increased substantially for men (from 4.6% to 10%).

Another implication is that the expected wage a worker earns later in his career potentially decreases when a positive discrimination policy favoring the worker is introduced (Proposition 4). As a consequence, it is possible that such policies do not diminish average "wage gaps" between different groups of workers. Support for this implication is provided by several studies. Smith and Welch (1984) study wage gaps between white and black men in the US between 1967 and 1980. They show that the wage gap became narrower in the first years of their analysis (until 1972), but that it stabilized after effective monitoring systems for affirmative action were established. The study by Leonard (1996) focuses on the US labor market in the 1980s and comes to a similar conclusion: "Between 1980 and 1990 the average wage gap decreased dramatically for women but increased for every other group." The most striking evidence is provided by Burger and Jafta (2012), who study the impact of affirmative action programs on employment and wages in South Africa. They show that the average wage gap between white and black workers increased substantially after the affirmative action policies were put in place. As argued before, these at first sight perverse effects are consistent with the results of our model.

When we investigated the policy's long-run effects, we showed that the policy also affects decisions in the first period and we discuss the corresponding empirical implications next. A first observation is that firms may change the composition of their workforce in the first period, in order to cope better with the policy's restrictions in the future (Proposition 5). In particular, firms make hiring decisions such that the ratio of the measures of hired workers from the two groups leads to optimal promotion standards maximizing expected average output per worker. This could be understood as a type of trickle-down effect, meaning that the policy not only affects which (older) workers firms promote to the high-level job, but also which (younger) workers they hire into the low-level job. Bertrand et al. (2014) do not observe such trickle-down effects in their study of Norwegian companies, but, as argued before, it is conceivable that they studied a transitional period in which firms are still adapting to the policy's introduction. It would thus be interesting to repeat the analysis of Bertrand et al. (2014) in a few years to study the policy's long-run effects.

Another finding of our model is that, after the introduction of the policy, a worker's total expected wage could be different from the worker's expected output (Proposition 5). Hiring decisions are interdependent with the consequence that the zero-profit condition holds for the entirety of workers a firm employs, but not necessarily on the individual level. More specifically, in the considered equilibrium all workers receive exactly the same expected total wage regardless of whether they emanate from group A or B. If there are productivity differences between the groups, an implication is that workers of one group are overpaid in the sense that they receive an expected wage exceeding their expected output, whereas workers from the other group are underpaid. In other words, an empirical implication of the model is that positive discrimination programs lead to systematic differences between wages and productivity for workers emanating from different groups. We are not aware of any study investigating such differences, but believe that this would be an interesting avenue for future research.

Finally, our model predicts that all external job moves are lateral in the sense that workers stay at the same job level when moving to an external firm. This prediction is independent of whether or not the positive discrimination policy is in place. When the policy is in place, however, the result hinges on the assumption of symmetric firms, which implies that firms manage to hire the very same measures of workers that they lose to external firms (conditional on group affiliation and job assignment), so that the newly-hired workers can simply replace the leaving workers. If we extended the model (e.g., Proposition 2) to allow for asymmetric firms, the results may be different. If firms differed in size, for instance, a smaller firm may hire more workers than it loses to the external firms. In this case, it is conceivable that, due to the quota constraint, the firm cannot assign all workers to the high-level job who were promoted by their initial employers. Two reactions by the firm are possible. To meet the quota constraint, it may either promote some workers from group A to the high-level job who were assigned to the low-level job by their initial employer, or it may demote some workers from group B to the low-level job who were initially assigned to the high-level job. The model could thus be modified such that the policy would lead to some promotions across firms for workers favored by the policy and demotions across firms for the workers disadvantaged by the policy. Again, to the best of our knowledge, there exists no study investigating this prediction.

6 Conclusion

In this paper, we consider a model in which promotions are used as signals of worker ability, and we examine the impact of a positive discrimination policy. The policy lowers the promotion standard for the workers who are discriminated against. In the short run, this is beneficial for the workers in the middle of the ability distribution, because these workers are promoted if and only if the policy is in place. In contrast, workers of either high or low ability generally suffer from the policy. This is because the policy does not change the promotion decision for these workers, but rather weakens the positive signal of being promoted and strengthens the negative signal of not being promoted. These findings imply that policies aimed at "leveling the playing field" are not always as beneficial as they may appear. If workers succeed in spite of many obstacles, the labor market learns a great deal about their characteristics, so it can reward the workers generously on this basis.[29]

In the long run, the policy benefits all targeted workers. The policy increases the wage that these workers earn early in their career, and this wage increase more than outweighs the possible future wage reduction. We derive several empirical implications from our model and we discuss empirical studies that are supportive of some of our predictions. We also offer new implications that researchers could take to the data.

[29] A related argument is advanced in Krishnamurthy and Edlin (2014). In a study of college admission rules, they find that stereotypes against a disadvantaged group of students can only be eliminated if these students face higher admission standards. Formally, they assume that the ability distributions of the disadvantaged and the advantaged students satisfy the monotone likelihood ratio property. This assumption implies that the expected ability of admitted students can only be equalized across groups when the students of the disadvantaged groups have to meet a higher standard in order to be admitted to college.

Appendix

Proofs:

Proof of Lemma 1. a) If the first-period employer managed to hire a worker again in the second period, the worker's second-period output would be at least

$$(1+S)(c_1 + d_1 \underline{a}_l).$$

In contrast, if one of the external firms managed to hire a worker, the worker's second-period output would be at most

$$c_2 + d_2 \overline{a}_l \,(= \max\{c_1 + d_1 \overline{a}_l, c_2 + d_2 \overline{a}_l\}).$$

Suppose that

$$(1+S)(c_1 + d_1 \underline{a}_l) > c_2 + d_2 \overline{a}_l$$
$$\Leftrightarrow S > (c_2 + d_2 \overline{a}_l)/(c_1 + d_1 \underline{a}_l) - 1.$$

It is then easy to understand that no equilibrium exists in which one of the external firms manages to hire a worker away from the first-period employer. To show this, denote the highest offer that worker j receives from one of the external firms by w_{j2}. According to Assumption 1b), the corresponding firm makes the same wage offer to all other workers from the same group and assigned to the same job. Notice that this set of workers has positive measure. This follows directly from the assumptions that the incumbent firm initially hired a group of workers with positive measure and that $\overline{a}_l$ is so high that there is at least one ability level such that the firm wants to promote workers of that ability to the high-level job at the end of the first period.

By Assumption 1c) we have $w_{j2} < (1+S)(c_1 + d_1 \underline{a}_l)$. Hence, the first-period employer always reacts by matching w_{j2} and thus retain all the workers. This proves part a) of the lemma with

$$S_1 := (c_2 + d_2 \overline{a}_l)/(c_1 + d_1 \underline{a}_l) - 1 > 0.$$

b) Suppose that $S > S_1$. Then the expected ability of workers that are actually switching firms is equal to the overall expected ability of workers (conditional on job assignment at the end of period 1). This is because workers are never successfully hired away (as shown in part a)), but a small fraction γ of workers leaves the first-period employer for reasons that are unrelated to ability and job assignment.

Consider an arbitrary worker j of group l who is assigned to job k. According to part b) of Assumption 1, the highest wage offer of an external firm for this worker is also made to all other workers of group l who are assigned to the same job. The workers' first-period employer therefore has an incentive to exactly match this highest wage offer from the external firms, meaning that it remains to determine this offer, which we denote by w_{l2}. Suppose that $w_{l2} < \max\left\{c_1 + d_1 E\left[a_{jl}\left|a_{jl} \in \tilde{A}_{kl}\right.\right], c_2 + d_2 E\left[a_{jl}\left|a_{jl} \in \tilde{A}_{kl}\right.\right]\right\} =: Z$. In this case, there is at least one external firm that gains by deviating and offering a wage from the interval (w_{l2}, Z). Thus, in equilibrium we never observe $w_{l2} < Z$. According to part c) of Assumption 1 none of the external firms offers a wage above Z so that we never observe $w_{l2} > Z$. Consequently, in equilibrium we have $w_{l2} = Z$.

Since the preceding arguments were made for an arbitrary worker from group l, the corresponding wage formula is the same for all workers from this group. ■

Proof of Proposition 1. a) Recall that the promotion rule is the same for all workers from group l. Consider arbitrary $\breve{a}_1 \in A_{1l}$ and $\breve{a}_2 \in A_{2l}$, meaning that the firm does not want to promote a worker from group l when the worker has ability $\breve{a}_1$, but that it wishes to promote a worker with ability $\breve{a}_2$. In particular, when all workers have the same ability, all workers are promoted if the ability level is $\breve{a}_2$, whereas none of the workers is promoted if the ability level is $\breve{a}_1$. This implies the following set of inequalities:

$$\begin{aligned}(1+S)(c_2 + d_2\breve{a}_2) - w_{l2}(2) &\geq (1+S)(c_1 + d_1\breve{a}_2) - w_{l2}(1) \\ (1+S)(c_2 + d_2\breve{a}_1) - w_{l2}(2) &< (1+S)(c_1 + d_1\breve{a}_1) - w_{l2}(1).\end{aligned} \tag{A1}$$

Rearranging the two conditions leads to

$$\begin{aligned}(1+S)(c_2 - c_1 + (d_2 - d_1)\breve{a}_2) \quad &\geq \quad w_{l2}(2) - w_{l2}(1) \\ &> \quad (1+S)(c_2 - c_1 + (d_2 - d_1)\breve{a}_1),\end{aligned} \tag{A2}$$

which immediately implies $\breve{a}_2 > \breve{a}_1$. Hence, because we assumed that $\overline{a}_l$ is so high that there is at least one ability level such that the firm wants to promote workers of that ability, the two sets A_{1l} and A_{2l} must be of the form $A_{1l} = [\underline{a}_l, a_l^p)$ and $A_{2l} = [a_l^p, \overline{a}_l]$, with $a_l^p \in (\underline{a}_l, \overline{a}_l)$.

b) In equilibrium the external firms correctly anticipate the firm's promotion rule as specified in a), i.e., $\tilde{A}_{kl} = A_{kl}$ for $k = 1, 2$. Thus, workers' second-period wages (conditional on job assignment) can be stated as

$$w_{l2}(1) = \max\left\{c_1 + d_1 E\left[a_{jl}|a_{jl} < a_l^p\right], c_2 + d_2 E\left[a_{jl}|a_{jl} < a_l^p\right]\right\}$$

and

$$w_{l2}(2) = \max\{c_1 + d_1 E[a_{jl}|a_{jl} \geq a_l^p], c_2 + d_2 E[a_{jl}|a_{jl} \geq a_l^p]\} > w_{l2}(1).$$

Since $E[a_{jl}|a_{jl} < a_l^p] \leq E[a_{jl}] < a^e$, it follows that

$$c_1 + d_1 E[a_{jl}|a_{jl} < a_l^p] > c_2 + d_2 E[a_{jl}|a_{jl} < a_l^p]$$

implying

$$w_{l2}(1) = c_1 + d_1 E[a_{jl}|a_{jl} < a_l^p].$$

Because $a_l^p \in A_{2l}$ and $w_{l2}(2) - w_{l2}(1) > 0$, inequality (A2) leads to

$$(1+S)(c_2 - c_1 + (d_2 - d_1) a_l^p) > 0 \quad \Rightarrow \quad c_2 + d_2 a_{jl} > c_1 + d_1 a_{jl} \quad \text{for all} \quad a_{jl} \geq a_l^p.$$

This immediately implies $c_2 + d_2 E[a_{jl}|a_{jl} \geq a_l^p] \geq c_1 + d_1 E[a_{jl}|a_{jl} \geq a_l^p]$ so that

$$w_{l2}(2) = c_2 + d_2 E[a_{jl}|a_{jl} \geq a_l^p].$$

c) Because $\Delta\pi_l^p(a_{jl}, a_l^p)$ corresponds to $\pi_{ijl}(2) - \pi_{ijl}(1)$, this function obviously characterizes the additional second-period profit from promotion. From (A1) it follows that the promotion threshold fulfills $\Delta\pi_l^p(a_l^p, a_l^p) = 0$.

d) Because

$$\begin{aligned}
\Delta\pi_l^p(a^e, a_l^p) &= (1-\gamma)[(1+S)(c_2 - c_1 + (d_2 - d_1) a^e) - (w_{l2}(2, a_l^p) - w_{l2}(1, a_l^p))] \\
&= (1-\gamma)[w_{l2}(1, a_l^p) - w_{l2}(2, a_l^p)] < 0,
\end{aligned}$$

$\Delta\pi_l^p(a_l^p, a_l^p) = 0$, and $\Delta\pi_l^p(a_{jl}, a_l^p)$ is strictly increasing in a_{jl} it follows that $a_l^p > a^e$. ■

Proof of Proposition 2. First, we note that, analogous to our argumentation in the basic model, if S is sufficiently high, no equilibrium exists in which one of the external firms ever manages to hire a worker away from the first-period employer (unless the worker switches firms for exogenous reasons). We can always choose S_2 such that this is fulfilled. We use backward induction to determine optimal behavior. In the proofs of parts b) to d), we assume that $E[a_{jA}|a_{jA} \geq a_A^q] \geq a^e$. In the proof of part a) we will show that, if $\alpha < \alpha_1$, firms always choose the promotion standard for group A such that this condition is fulfilled.

Proof of d): We begin by showing that the firms have no incentive to deviate from the job assignment of newly-hired external workers in $\tau = 2$ as described in part d) if firms and workers follow the proposed behavior up to this decision. Each firm originally hired a continuum of workers from group A or group B of measure n_A/N or n_B/N, respectively. The individual abilities a_{jl} of each worker j of group $l \in \{A, B\}$ are assumed to be independently and identically distributed over $[\underline{a}_l, \overline{a}_l]$. Let I denote the indicator function of set $[0, \infty)$, i.e., $I(x) = 1$ if $x \geq 0$ and $I(x) = 0$ otherwise. Against the background of the law of large numbers in the case of a continuum of workers[30]

$$E[I(a_{jl} - x)] = prob(a_{jl} > x) = 1 - F_l(x)$$

corresponds to the proportion of employees in some arbitrary firm emanating from group l with ability greater than x and

$$\frac{n_l}{N} \cdot (1 - F_l(x))$$

stands for the corresponding measure of employees.

It follows that each firm loses a continuum of workers from group l that were assigned to job 1 of measure $\gamma F_l\left(a_l^q\right) \frac{n_l}{N}$ and to job 2 of measure $\gamma\left(1 - F_l\left(a_l^q\right)\right) \frac{n_l}{N}$ in period 2. Because all firms were initially of equal size, each firm simultaneously manages to hire the very same measures of workers from the same group and assigned to the same job. In equilibrium, all firms can correctly anticipate the chosen promotion standards a_l^q. Because $E\left[a_{jl} | a_{jl} \geq a_l^q\right] \geq a^e$ for both l (for group A this was assumed before, for group B the condition holds since $a_B^q > a^e$), each firm would find it optimal to assign all newly-hired workers that were originally assigned to job 2 to the high-level job as well. Since the quota constraint was originally met and the newly-hired workers simply replace the leaving workers, such assignment is always possible without violating the quota constraint. Similarly, because $E\left[a_{jl} | a_{jl} < a_l^q\right] < E\left[a_{jl}\right] < a^e$ for both l, each firm always assigns all newly-hired workers that were originally assigned to job 1 to the low-level job as well. To sum up, when firms follow the equilibrium from the proposition up to the point at which they decide about the assignment to jobs of the newly-hired workers, none of the firms wishes to deviate from the assignment described in the proposition.

Proof of c): Obvious and therefore omitted.

Proof of b): The next decision relates to firms' period-2 wage offers. First, it is easy to see that a firm's wage offers to the own employees are always optimal since matching the external firms' offers represents the cheapest way to retain the workers. Consider now the wage offers to the external firms' workers. Again suppose that all firms follow the behavior

[30] For the law of large numbers in the continuous case see Al-Najjar (2004).

described in the proposition up to this decision. If firms continued to stick to the proposed equilibrium behavior, each firm would earn zero expected profit from the newly-hired workers. The reason is that each worker that was originally assigned to job k will be assigned to the very same job when switching firms (as explained in the proof of d)), implying that the period-2 wage offers specified in the proposition are equal to the workers' expected period-2 outputs. When lowering some or all of the wage offers, a firm would never have a chance at hiring the workers whose offers were lowered, meaning that the measure of the hired workers would decline. In addition, the firm might have to reassign some of the remaining workers to the job at which they are least productive in case the quota constraint were violated (e.g., if only the wage offers to workers from group A were lowered). Hence, such deviation from the equilibrium can never be profitable. Increasing the wage offers to some or all of the workers would not be profitable either since the firm would then pay those workers a wage above their expected output and, thus, suffer a loss when they indeed choose to switch firms. This could only be profitable when this would relax the quota constraint, making it easier to assign other workers to the job at which they are most productive. As the newly-hired workers are already assigned to the jobs at which they are most productive and the existing workers cannot be reassigned, this is not the case. Summing up, none of the firms has an incentive to deviate from the period-2 wage offers when all firms stick to the behavior described in the proposition until period-2 wage offers are made.

Proof of a): Finally, we turn to firms' promotion decisions at the end of $\tau = 1$. Let i be an arbitrary firm. If the firm applies the optimal promotion standards a_A^p and a_B^p, respectively,

$$q_{Ai}^p(a_A^p, a_B^p) := \frac{n_{Ai} \cdot (1 - F_A(a_A^p))}{n_{Ai} \cdot (1 - F_A(a_A^p)) + n_{Bi} \cdot (1 - F_B(a_B^p))}$$

equals the promotion share of group-A-workers. As long as $q_{Ai}^p(a_A^p, a_B^p) \geq \alpha$ the introduction of quota α has no influence on the optimal promotion behavior of the firm.

However, in the case $q_{Ai}^p(a_A^p, a_B^p) < \alpha$ the proportion $q_{Ai}^p(a_A^p, a_B^p)$ needs to be increased. Because $\partial q_{A_i}^p / \partial a_A^p < 0$ and $\partial q_{A_i}^p / \partial a_B^p > 0$, new promotion standards $a_A^q \leq a_A^p$ or $a_B^q \geq a_B^p$ need to be chosen such that $q_{Ai}^p(a_A^q, a_B^q) \geq \alpha$. Furthermore, we assumed $\bar{a}_l$ ($l \in \{A, B\}$) to be so high that there is at least one ability level such that the firm wants to promote workers of that ability. This immediately leads to $a_l^q < \bar{a}_l \Leftrightarrow F_l(a_l^q) < 1$. In equilibrium, the standards a_A^q and a_B^q always have to fulfill the conditions $\Delta\pi_A^p(a_A^q, a_A^q) \leq 0$ and $\Delta\pi_B^p(a_B^q, a_B^q) \geq 0$

because firm i could always deviate from the proposed promotion standards by lowering the standard for group A or increasing the standard for group B without violating the quota constraint. Such deviations are not profitable for the firm if the two inequalities are met.

In addition, we have

$$\begin{aligned}\Delta\pi_l^p\left(a^e, a^e\right) &= (1-\gamma)[(1+S)(c_2 - c_1 + (d_2 - d_1)\, a^e) - (w_{l2}(2, a^e) - w_{l2}(1, a^e))] \\ &= (1-\gamma)[w_{l2}(1, a^e) - w_{l2}(2, a^e)] < 0,\end{aligned}$$

and

$$\begin{aligned}\Delta\pi_l^p\left(\bar{a}_l, \bar{a}_l\right) &= (1-\gamma)[(1+S)(c_2 - c_1 + (d_2 - d_1)\, \bar{a}_l) - (w_{l2}(2, \bar{a}_l) - w_{l2}(1, \bar{a}_l))] \\ &= (1-\gamma)[(1+S)(c_2 - c_1 + (d_2 - d_1)\, \bar{a}_l) - (c_2 + d_2\bar{a}_l - c_1 - d_1 E\left[a_{jl}\right])] > 0\end{aligned}$$

if $S_2(< S)$ is sufficiently high. Because $a^e < a_l^p < \bar{a}_l$, $\Delta\pi_l^p(a, a)$ is continuous in a, and we assumed (throughout Section 4) the solution for a_l^p to the condition $\Delta\pi_l^p(a_l^p, a_l^p) = 0$ to be unique, it immediately follows that $\Delta\pi_l^p(a, a) < 0$ for all $a < a_l^p$ and $\Delta\pi_l^p(a, a) > 0$ for all $a > a_l^p$. Under consideration of $\Delta\pi_A^p(a_A^q, a_A^q) \leq 0$ and $\Delta\pi_B^p(a_B^q, a_B^q) \geq 0$ it follows $a_A^q \leq a_A^p$ and $a_B^q \geq a_B^p$.

As mentioned in the introduction of the proof, it remains to show $E\left[a_{jA}|a_{jA} \geq a_A^q\right] \geq a^e$. This is immediately fulfilled if α is only slightly above $q_{Ai}^p(a_A^p, a_B^p)$ (i.e., below some threshold α_1) and consequently, the promotion standards a_A^q and a_B^q will not differ strongly from a_A^p and a_B^p.

Finally, since all firms hired a continuum of workers from groups A and B of measures $\frac{n_A}{N}$ and $\frac{n_B}{N}$, there is always an equilibrium in which they all find it optimal change the promotion standards in the same way. ■

Proof of Proposition 3. a) The statement immediately follows from the wage formula in part b) of Proposition 2 because $c_1 + d_1 E\left[a_{jl}|a_{jl} < a\right]$ as well as $c_2 + d_2 E\left[a_{jl}|a_{jl} \geq a\right]$ are increasing functions in a.

b) Case 1: $a_{jA} \in [\underline{a}_A, a_A^q)$

In this case the worker is neither promoted if the promotion standard is a_A^p nor if the promotion standard is a_A^q. Consequently, the period-2 wage difference amounts to

$$w_{A2}(1, a_A^q) - w_{A2}(1, a_A^p) = d_1\left(E\left[a_{jA}\left|a_{jA} < a_A^q\right.\right] - E\left[a_{jA}\left|a_{jA} < a_A^p\right.\right]\right) < 0.$$

The latter inequality again results from the fact that $E[a_{jl}|a_{jl} < a]$ is increasing in a.

Case 2: $a_{jA} \in [a_A^p, \overline{a}_A]$

In this case the worker is promoted for both promotion standards a_A^p and a_A^q and it can be analogously shown that $w_{A2}(2, a_A^q) - w_{A2}(2, a_A^p) < 0$.

Case 3: $a_{jA} \in [a_A^q, a_A^p)$

In this case the worker is not promoted if the promotion standard is a_A^p but is promoted if the promotion standard is a_A^q. The resulting wage difference corresponds to

$$\begin{aligned} & w_{A2}(2, a_A^q) - w_{A2}(1, a_A^p) \\ = \; & c_2 - c_1 + d_1 \left(E\left[a_{jA} \left| a_{jA} \geq a_A^q \right.\right] - E\left[a_{jA} \left| a_{jA} < a_A^p \right.\right]\right) \\ & +(d_2 - d_1) E\left[a_{jA} \left| a_{jA} \geq a_A^q \right.\right] \\ > \; & c_2 - c_1 + (d_2 - d_1) E\left[a_{jA} \left| a_{jA} \geq a_A^q \right.\right] \geq 0. \end{aligned}$$

The second-to-last inequality results because

$$E\left[a_{jA} \left| a_{jA} \geq a_A^q \right.\right] > E(a_j) > E\left[a_{jA} \left| a_{jA} < a_A^p \right.\right]$$

and the last inequality because $E[a_{jA}|a_{jA} \geq a_A^q] \geq a^e$. ■

Proof of Proposition 4. a) W_{A2} corresponds to

$$\begin{aligned} W_{A2}(a) \; = \; & F_A(a) \cdot (c_1 + d_1 \cdot E[a_{jA}|a_{jA} < a]) \\ & + (1 - F_A(a)) \cdot (c_2 + d_2 \cdot E[a_{jA}|a_{jA} \geq a]) \end{aligned} \tag{A3}$$

and can be restated as

$$\begin{aligned} W_{A2}(a) &= c_2 + (c_1 - c_2) F_A(a) + d_1 E(a_{jA}) + (d_2 - d_1)(1 - F_A(a)) \cdot E[a_{jA}|a_{jA} \geq a] \\ &= c_2 + (c_1 - c_2) F_A(a) + d_1 E(a_{jA}) + (d_2 - d_1) \int_a^{\overline{a}_A} a_{jA} f_A(a_{jA}) da_{jA}. \end{aligned}$$

The derivative

$$W'_{A2}(a) = (c_1 - c_2 - (d_2 - d_1) \cdot a) \cdot f_A(a) \tag{A4}$$

leads to the following necessary condition of an extremum:

$$W'_{A2}(a^*) = 0 \quad \Leftrightarrow \quad a^* = \frac{c_1 - c_2}{d_2 - d_1} = a^e.$$

Furthermore,

$$W''_{A2}(a) = -(d_2 - d_1) f_A(a) + (c_1 - c_2 - (d_2 - d_1) \cdot a) \cdot f'_A(a).$$

Because $(c_1 - c_2) - (d_2 - d_1)a^e = 0$, it immediately results $W''_{A2}(a^e) < 0$ which implies a local maximum of W_{A2} at a^e. Because a^e is the only zero of W'_{A2} in $(\hat{a}, \overline{a}_A)$ and W_{A2} is continuous over $[\hat{a}, \overline{a}_A]$, a^e is the global maximum of W_{A2}.

b) Notice that

$$
\begin{aligned}
& c_2 + d_2 \cdot E\left[a_{jA} | a_{jA} \geq \hat{a}\right] = c_2 + d_2 \cdot a^e \\
= \;& c_1 + d_1 \cdot a^e = c_1 + d_1 \cdot E\left[a_{jA} | a_{jA} \geq \hat{a}\right].
\end{aligned}
$$

By the law of total expectation, it then follows that

$$
\begin{aligned}
W_{A2}(\hat{a}) &= F_A(\hat{a}) \cdot \left(c_1 + d_1 \cdot E\left[a_{jA} | a_{jA} < \hat{a}\right]\right) \\
&\quad + \left(1 - F_A(\hat{a})\right) \cdot \left(c_2 + d_2 \cdot E\left[a_{jA} | a_{jA} \geq \hat{a}\right]\right) \\
&= F_A(\hat{a}) \cdot \left(c_1 + d_1 \cdot E\left[a_{jA} | a_{jA} < \hat{a}\right]\right) \\
&\quad + \left(1 - F_A(\hat{a})\right) \cdot \left(c_1 + d_1 \cdot E\left[a_{jA} | a_{jA} \geq \hat{a}\right]\right) \\
&= c_1 + d_1 \cdot E\left[a_{jA}\right].
\end{aligned}
$$

Since $W'_{A2}(a) < 0$ for all $a \in (a^e, \overline{a}_A) \supset (a^p_A, \overline{a}_A)$, we have

$$
W_{A2}(\hat{a}) = c_1 + d_1 \cdot E\left[a_{jA}\right] = W_{A2}(\overline{a}_A) < W_{A2}(a^p_A).
$$

Because, moreover, W_{A2} is continuous, part b) is proven. ■

Proof of Proposition 5. First, we note that there is the same threshold value S_2 as in Proposition 2 such that, if $S > S_2$, no equilibrium exists in which one of the external firms ever manages to hire a worker away from the first-period employer (unless the worker switches firms for exogenous reasons). We use backward induction to determine optimal behavior.

Proof of f): The proof is completely analogous to the proof of part d) of Proposition 2.

Proof of e): Obvious and therefore omitted.

Proof of d): For $\hat{N} \geq 3$, the proof is analogous to the proof of part b) of Proposition 2 with the exception that we still have to show that the $N - \hat{N}$ inactive firms do not want to deviate by offering a wage equal to or higher than $w_{l2}(k, a_l^{p*})$. This is straightforward. In the proposed equilibrium, the $N - \hat{N}$ inactive firms do not manage to hire any workers and, therefore, earn zero profit. By making a wage offer to a worker of at least $w_{l2}(k, a_l^{p*})$, the firms would offer the worker a wage at least equal to the worker's expected output when the worker is assigned efficiently to a job given the inactive firms' information about the worker's ability. Hence, such a deviation can never be profitable.

Proof of c): Next, we turn to firms' promotion decisions at the end of $\tau = 1$. Suppose again that up to these decisions, all firms behaved as described in the proposition, meaning that each firm hired a continuum of workers of measure $\hat{n}_A$ from group A and of measure $\hat{n}_B$ from group B. When deciding about which workers to promote, we already explained in the main model that each firm finds it optimal to follow a cutoff-rule, promoting only workers with sufficiently high ability. The only thing that we therefore have to check is whether a firm ever has an incentive to deviate from the cutoff-values a_A^{p*} and a_B^{p*} defined in the proposition. Notice first that by deviating from any of the standards, the firm could not affect the external firms' wage offers for the workers the firm currently employs. This is because the external firms cannot observe the actual choice of promotion standard, and hence the deviation, implying that their wage offers do not respond to a change in the promotion standards. Furthermore, deviating from any of the standards does not improve the assignment of newly-hired workers in $\tau = 2$ who are efficiently assigned given the firm's information about their abilities.

Start by assuming that $q_{Ai}^p(a_A^{p*}, a_B^{p*}) > \alpha$, i.e., that the quota constraint is slack when firm i chooses the proposed promotion standards. Then a firm could (marginally) change the promotion thresholds, thereby promoting relatively fewer or more workers, without violating the quota constraint. The firm does not want to change promotion standards because the condition

$$\Delta\pi_A^p(a_A^{p*}, a_A^{p*}) = 0 = \Delta\pi_B^p(a_B^{p*}, a_B^{p*})$$

is satisfied.

Suppose now that the quota constraint is binding, i.e., $q_A^p(a_A^{p*}, a_B^{p*}) = \alpha$. A firm could always deviate from the proposed promotion standards by lowering the standard for group A or increasing the standard for group B without violating the quota constraint. It does not want to do so because

$$\Delta\pi_A^p(a_A^{p*}, a_A^{p*}) \leq 0 \quad \text{and} \quad \Delta\pi_B^p(a_B^{p*}, a_B^{p*}) \geq 0$$

hold. In contrast, if a firm would deviate by increasing the standard for group A, the standard for group B would have to be increased as well in order to fulfill

$$q_{Ai}^p(a_A^p, a_B^p) = \alpha \quad \Leftrightarrow \quad a_B^p(a_A^p) := F_B^{-1}\left(1 - \frac{(1-\alpha)\cdot n_{A_i}}{\alpha \cdot n_{B_i}} \cdot (1 - F_A(a_A^p))\right).$$

If we define

$$\pi_l^p(a_l^p, a_l^{p*}) = (1-\gamma) \cdot n_{l_i} \cdot [(1+S)Y_{l2}(a_l^p) - (F_l(a_l^p) w_{l2}(1, (a_l^{p*})) + (1 - F_l(a_l^p))\, w_{l2}(2, (a_l^{p*})))]$$

as the expected second-period profit of group l, a firm would only deviate by increasing the standard a_A^p for group A if

$$\frac{d}{da_A^p} \left(\pi_A^p(a_A^p, a_A^{p*}) + \pi_B^p(a_B^p(a_A^p), a_B^{p*})\right) > 0,$$

i.e., if the total expected second-period profit would be increased. The latter inequality is equivalent to[31]

$$-\alpha \Delta\pi_A^p(a_A^p, a_A^{p*}) - (1-\alpha)\, \Delta\pi_B^p(a_B^p(a_A^p), a_B^{p*}) > 0.$$

Because in equilibrium we have $a_A^p = a_A^{p*}$ and $a_B^p(a_A^p) = a_B^p(a_A^{p*}) = a_B^{p*}$, the latter inequality contradicts the condition $\alpha \Delta\pi_A^p(a_A^{p*}, a_A^{p*}) + (1-\alpha)\, \Delta\pi_B^p(a_B^{p*}, a_B^{p*}) = 0$ of the auxiliary problem (AP_i). Analogously, we can show that the firm does not wish to deviate by lowering the standard for group B and adjusting the standard for A such that the quota constraint remains fulfilled. Such a deviation would require $\alpha \Delta\pi_A^p(a_A^{p*}, a_A^{p*}) + (1-\alpha)\, \Delta\pi_B^p(a_B^{p*}, a_B^{p*}) > 0$, again contradicting the condition $\alpha \Delta\pi_A^p(a_A^{p*}, a_A^{p*}) + (1-\alpha)\, \Delta\pi_B^p(a_B^{p*}, a_B^{p*}) = 0$ of the auxiliary problem (AP_i). Summing up, we have shown that firms do not wish to deviate from the proposed promotion standards as described by the constraints of auxiliary problem (AP_i).

Proof of b): Given that the $\hat{N}$ firms offer the same wages in each period and given that they choose the same promotion rules (which the workers correctly anticipate), the workers assess the $\hat{N}$ firms identically. Because only wages are the basis for workers' choice, they accept each of the $\hat{N}$ offers with the same probability $\frac{1}{\hat{N}}$ at the beginning of $\tau = 1$.

Proof of a): It remains to be shown that none of the firms wishes to deviate from the period-1 wage offers. We notice first that, in the proposed equilibrium, each firm earns zero expected profit (regardless of whether or not it belongs to the $\hat{N}$ firms that manage to hire workers). To see this, we consider the expected wage costs of the initially hired workers which

[31] The proof of this equivalence is available from the authors upon request.

in the proposed equilibrium correspond to

$$
\begin{aligned}
& \hat{n}_A \left(w_{A1}^* + (1-\gamma) \left[F_A \left(a_A^{p*} \right) w_{A2} \left(1, a_A^{p*} \right) + \left(1 - F_A \left(a_A^{p*} \right) \right) w_{A2} \left(2, a_A^{p*} \right) \right] \right) \\
& + \hat{n}_B \left(w_{B1}^* + (1-\gamma) \left[F_B \left(a_B^{p*} \right) w_{B2} \left(1, a_B^{p*} \right) + \left(1 - F_B \left(a_B^{p*} \right) \right) w_{B2} \left(2, a_B^{p*} \right) \right] \right) \\
= \; & \hat{n}_A \left(w_{A1}^* + (1-\gamma) Y_{A2} \left(a_A^{p*} \right) \right) + \hat{n}_B \left(w_{B1}^* + (1-\gamma) Y_{B2} \left(a_B^{p*} \right) \right) \\
\underset{a)}{=} \; & \frac{\hat{n}_A \hat{n}_B}{\hat{n}_A + \hat{n}_B} \left[Y_{B1} + \frac{\hat{n}_A}{\hat{n}_B} Y_{A1} + (1 + (1-\gamma) S) Y_{B2} \left(a_B^{p*} \right) + \left(\frac{\hat{n}_A}{\hat{n}_B} (1-\gamma) S - 1 \right) Y_{A2} \left(a_A^{p*} \right) \right] \\
& + \hat{n}_A (1-\gamma) Y_{A2} \left(a_A^{p*} \right) \\
& + \frac{\hat{n}_B \hat{n}_A}{\hat{n}_A + \hat{n}_B} \left[Y_{A1} + \frac{\hat{n}_B}{\hat{n}_A} Y_{B1} + (1 + (1-\gamma) S) Y_{A2} \left(a_A^{p*} \right) + \left(\frac{\hat{n}_B}{\hat{n}_A} (1-\gamma) S - 1 \right) Y_{B2} \left(a_B^{p*} \right) \right] \\
& + \hat{n}_B (1-\gamma) Y_{B2} \left(a_B^{p*} \right) \\
= \; & \hat{n}_A Y_{A1} + \hat{n}_B Y_{B1} + \hat{n}_A (1-\gamma)(1+S) Y_{A2} \left(a_A^{p*} \right) + \hat{n}_B (1-\gamma)(1+S) Y_{B2} \left(a_B^{p*} \right).
\end{aligned}
$$

We observe that the firm's expected wage cost equals the expected output and expected profits are zero.

We therefore have to show that it is not possible for a firm to earn strictly positive profit by choosing period-1 wage offers different from those specified in the proposition.

Suppose that some firm i deviates from the period-1 wages stated in the proposition and that it attracts a continuum of workers from group A of measure n_{Ai} and from group B of measure n_{Bi}. The firm then chooses promotion standards a_{Ai}^p and a_{Bi}^p that the other firms and all workers correctly anticipate (since firms and workers are able to observe n_{Ai} and n_{Bi}).

In the equilibrium considered in the proposition, a worker from group l receives expected total payoff equal to

$$
w_{l1}^* + F_l \left(a_l^{p*} \right) w_{l2} \left(1, a_l^{p*} \right) + \left(1 - F_l \left(a_l^{p*} \right) \right) w_{l2} \left(2, a_l^{p*} \right).
$$

To hire a worker from group l, firm i at least would have to match this payoff, meaning that the condition

$$
\begin{aligned}
& w_{l1i} + F_l \left(a_{li}^{p} \right) w_{l2} \left(1, a_{li}^{p} \right) + \left(1 - F_l \left(a_{li}^{p} \right) \right) w_{l2} \left(2, a_{li}^{p} \right) \\
\geq \; & w_{l1}^* + F_l \left(a_l^{p*} \right) w_{l2} \left(1, a_l^{p*} \right) + \left(1 - F_l \left(a_l^{p*} \right) \right) w_{l2} \left(2, a_l^{p*} \right)
\end{aligned}
$$

would have to be met, which defines a lower bound for firm i's period-1 wage offer w_{l1i}. This lower bound immediately leads to an upper bound for firm i's profit π_{li}^{dev} from hiring a continuum of workers from group l of measure n_{li} which is given by

$$\begin{aligned}
\pi_{li}^{dev} &\leq n_{li}Y_{l1} + n_{li}(1-\gamma)(1+S)Y_{l2}(a_{li}^{p}) \\
&\quad - n_{li}(1-\gamma)\left(F_l(a_{li}^{p})w_{l2}(1,a_{li}^{p}) + (1-F_l(a_{li}^{p}))w_{l2}(2,a_{li}^{p})\right) \\
&\quad - n_{li}\left(w_{l1}^{*} + F_l(a_l^{p*})w_{l2}(1,a_l^{p*}) + (1-F_l(a_l^{p*}))w_{l2}(2,a_l^{p*})\right) \\
&\quad + n_{li}\left(F_l(a_{li}^{p})w_{l2}(1,a_{li}^{p}) + (1-F_l(a_{li}^{p}))w_{l2}(2,a_{li}^{p})\right) \\
&= n_{li}Y_{l1} + n_{li}(1-\gamma)(1+S)Y_{l2}(a_{li}^{p}) \\
&\quad + n_{li}\gamma\left(F_l(a_{li}^{p})w_{l2}(1,a_{li}^{p}) + (1-F_l(a_{li}^{p}))w_{l2}(2,a_{li}^{p})\right) \\
&\quad - n_{li}\left(w_{l1}^{*} + F_l(a_l^{p*})w_{l2}(1,a_l^{p*}) + (1-F_l(a_l^{p*}))w_{l2}(2,a_l^{p*})\right) \\
&\leq n_{li}Y_{l1} + n_{li}(1-\gamma)(1+S)Y_{l2}(a_{li}^{p}) + n_{li}\gamma Y_{l2}(a_{li}^{p}) - n_{li}\left(w_{l1}^{*} + Y_{l2}(a_l^{p*})\right).
\end{aligned}$$

The latter inequality results from the fact that the expected wage offer of firm i in $\tau = 2$ does not exceed the period-2 output that the firm expects a worker to produce, i.e.,

$$\begin{aligned}
&F_l(a_{li}^{p})w_{l2}(1,a_{li}^{p}) + (1-F_l(a_{li}^{p}))w_{l2}(2,a_{li}^{p}) \\
&\leq F_l(a_{li}^{p})\left(c_1 + d_1E\left[a_{jl}|a_{jl} < a_{li}^{p}\right]\right) + (1-F_l(a_{li}^{p}))\left(c_2 + d_2E\left[a_{jl}|a_{jl} \geq a_{li}^{p}\right]\right) = Y_{l2}(a_{li}^{p}).
\end{aligned}$$

Under consideration of the first-period wage offers

$$\begin{aligned}
w_{A1}^{*} &= \frac{1}{1+q_A^{*}}\left(Y_{B1} + q_A^{*}Y_{A1} + (1+(1-\gamma)S)Y_{B2}(a_B^{p*}) + (q_A^{*}(1-\gamma)S - 1)Y_{A2}(a_A^{p*})\right), \\
w_{B1}^{*} &= \frac{1}{1+q_B^{*}}\left(Y_{A1} + q_B^{*}Y_{B1} + (1+(1-\gamma)S)Y_{A2}(a_A^{p*}) + (q_B^{*}(1-\gamma)S - 1)Y_{B2}(a_B^{p*})\right) \\
&= \frac{1}{1+q_A^{*}}\left(q_A^{*}Y_{A1} + Y_{B1} + q_A^{*}(1+(1-\gamma)S)Y_{A2}(a_A^{p*}) + ((1-\gamma)S - q_A^{*})Y_{B2}(a_B^{p*})\right)
\end{aligned}$$

of the other firms to workers of groups A and B, firm i's profit $\pi_i^{dev} = \pi_{Ai}^{dev} + \pi_{Bi}^{dev}$ from hiring a continuum of workers from groups A and B of measure n_{Ai} and n_{Bi} can be bounded from above as follows:

$$\begin{aligned}
\pi_i^{dev} &\leq n_{Ai}Y_{A1} + n_{Ai}(1-\gamma)(1+S)Y_{A2}(a_{Ai}^p) + n_{Ai}\gamma Y_{A2}(a_{Ai}^p) - n_{Ai}Y_{A2}(a_A^{p*}) \\
&\quad - n_{Ai}\frac{1}{1+q_A^*}\left(Y_{B1} + q_A^* Y_{A1} + (1+(1-\gamma)S)Y_{B2}(a_B^{p*}) + (q_A^*(1-\gamma)S - 1)Y_{A2}(a_A^{p*})\right) \\
&\quad + n_{Bi}Y_{B1} + n_{Bi}(1-\gamma)(1+S)Y_{B2}(a_{Bi}^p) + n_{Bi}\gamma Y_{B2}(a_{Bi}^p) - n_{Bi}Y_{B2}(a_B^{p*}) \\
&\quad - n_{Bi}\frac{1}{1+q_A^*}\left(q_A^* Y_{A1} + Y_{B1} + q_A^*(1+(1-\gamma)S)Y_{A2}(a_A^{p*}) + ((1-\gamma)S - q_A^*)Y_{B2}(a_B^{p*})\right) \\
&= n_{Ai}Y_{A1} + n_{Ai}(1+(1-\gamma)S)Y_{A2}(a_{Ai}^p) \\
&\quad - n_{Ai}\frac{1}{1+q_A^*}\left(Y_{B1} + q_A^* Y_{A1} + (1+(1-\gamma)S)Y_{B2}(a_B^{p*}) + q_A^*(1+(1-\gamma)S)Y_{A2}(a_A^{p*})\right) \\
&\quad + n_{Bi}Y_{B1} + n_{Bi}(1+(1-\gamma)S)Y_{B2}(a_{Bi}^p) \\
&\quad - n_{Bi}\frac{1}{1+q_A^*}\left(q_A^* Y_{A1} + Y_{B1} + q_A^*(1+(1-\gamma)S)Y_{A2}(a_A^{p*}) + (1+(1-\gamma)S)Y_{B2}(a_B^{p*})\right).
\end{aligned}$$

With $q_A^* = n_A^*/n_B^*$ the latter term is non-positive if and only if

$$\begin{aligned}
&\frac{n_{Ai}}{n_{Ai}+n_{Bi}}\left(Y_{A1} + (1+(1-\gamma)S)Y_{A2}(a_{Ai}^p)\right) + \frac{n_{Bi}}{n_{Ai}+n_{Bi}}\left(Y_{B1} + (1+(1-\gamma)S)Y_{B2}(a_{Bi}^p)\right) \\
\leq\ &\frac{n_A^*}{n_A^*+n_B^*}\left(Y_{A1} + (1+(1-\gamma)S)Y_{A2}(a_A^{p*})\right) + \frac{n_B^*}{n_A^*+n_B^*}\left(Y_{B1} + (1+(1-\gamma)S)Y_{B2}(a_B^{p*})\right).
\end{aligned}$$

Since q_B^*, a_A^{p*} and a_B^{p*} is the solution to problem (AP_i), this condition is always fulfilled and consequently, there is no incentive for firm i to deviate from the first-period wage offered in the proposed equilibrium. ■

Proof of Proposition 6. Recall that period-2 wages are independent of S regardless of whether or not the positive discrimination policy is in place. We will prove the proposition by showing that, under the imposed assumptions, the difference in period-1 wages for workers from group A when comparing the situations with and without the policy is increasing without bound in S.

Using Proposition 5 and formula (2), the corresponding difference in period-1 wages can be stated as

$$\begin{aligned}
w_{A1}^* - w_{A1} &= \frac{1}{1+q_A^*}\left(Y_{B1} + q_A^* Y_{A1} + (1+(1-\gamma)S)Y_{B2}(a_B^{p*}) + (q_A^*(1-\gamma)S - 1)Y_{A2}(a_A^{p*})\right) \\
&\quad - Y_{A1} - (1-\gamma)SY_{A2}(a_A^p).
\end{aligned}$$

We obtain

$$\frac{\partial(w_{A1}^* - w_{A1})}{\partial S} = (1-\gamma)\left(\frac{q_A^*}{1+q_A^*}Y_{A2}(a_A^{p*}) + \frac{1}{1+q_A^*}Y_{B2}(a_B^{p*}) - Y_{A2}(a_A^p)\right).$$

Notice that $n_{Ai} > 0$ and $n_{Bi} = 0$ (i.e. $q_{Ai} = 1$ and $q_{Bi} = 0$) would satisfy all the constraints from (AP_i); when firm i employs only workers from group A, the quota constraint is always fulfilled. Since the quota constraint does not restrict the firm's decisions, we have the same solution as in the basic model, i.e., the firm chooses promotion standard a_A^p, and the value of the objective function in (AP_i) is $Y_{A1} + (1 + (1-\gamma) S) Y_{A2} (a_A^p)$.

However, according to the proposition's conditions, any solution to (AP_i) has $q_A^* > 0$ and $q_B^* > 0$. This means that the corresponding value of the objective function must be higher than when choosing $n_{Bi} = 0$. Consequently, we must have

$$\begin{aligned} & \frac{q_A^*}{1+q_A^*} \left[Y_{A1} + (1+(1-\gamma) S) Y_{A2} (a_A^{p*})\right] + \frac{1}{1+q_A^*} \left[Y_{B1} + (1+(1-\gamma) S) Y_{B2} (a_B^{p*})\right] \\ > \; & Y_{A1} + (1+(1-\gamma) S) Y_{A2} (a_A^p) \\ \Leftrightarrow \; & (1+(1-\gamma) S) \left(\frac{q_A^*}{1+q_A^*} Y_{A2} (a_A^{p*}) + \frac{1}{1+q_A^*} Y_{B2} (a_B^{p*}) - Y_{A2} (a_A^p) \right) > \frac{1}{1+q_A^*} (Y_{A1} - Y_{B1}) . \end{aligned}$$

Hence, whenever $Y_{A1} \geq Y_{B1}$, we observe $\frac{q_A^*}{1+q_A^*} Y_{A2} (a_A^{p*}) + \frac{1}{1+q_A^*} Y_{B2} (a_B^{p*}) - Y_{A2} (a_A^p) > 0$, implying $\partial(w_{A1}^* - w_{A1})/\partial S$ to be positive and $w_{A1}^* - w_{A1}$ to be increasing without bound in S. ■

Proof of Proposition 7. Taking the results from Proposition 5 into account, a worker's expected total wage can be stated as

$$\begin{aligned} & w_{A1}^* + F_A (a_A^{p*}) w_{A2} (1, a_A^{p*}) + (1 - F_A (a_A^{p*})) w_{A2} (2, a_A^{p*}) \\ = \; & \frac{1}{1+q_A^*} \left(Y_{B1} + q_A^* Y_{A1} + (1+(1-\gamma) S) Y_{B2} (a_B^{p*}) + (q_A^* (1-\gamma) S - 1) Y_{A2} (a_A^{p*})\right) + Y_{A2} (a_A^{p*}) \\ = \; & \frac{1}{1+q_A^*} \left(Y_{B1} + q_A^* Y_{A1} + (1+(1-\gamma) S) Y_{B2} (a_B^{p*}) + q_A^* (1+(1-\gamma) S) Y_{A2} (a_A^{p*})\right) \\ = \; & \frac{q_A^*}{1+q_A^*} \left[Y_{A1} + (1+(1-\gamma) S) Y_{A2} (a_A^{p*})\right] + \frac{1}{1+q_A^*} \left[Y_{B1} + (1+(1-\gamma) S) Y_{B2} (a_B^{p*})\right] . \end{aligned}$$

In the basic model from Section 3 the corresponding wage is given by

$$Y_{A1} + (1+(1-\gamma) S) Y_{A2} (a_A^p) .$$

As shown in the proof of Proposition 6, when any solution to (AP_i) has $q_A^* > 0$ and $q_B^* > 0$, it follows that

$$\begin{aligned} & \frac{q_A^*}{1+q_A^*} \left[Y_{A1} + (1+(1-\gamma) S) Y_{A2} (a_A^{p*})\right] + \frac{1}{1+q_A^*} \left[Y_{B1} + (1+(1-\gamma) S) Y_{B2} (a_B^{p*})\right] \\ > \; & Y_{A1} + (1+(1-\gamma) S) Y_{A2} (a_A^p) . \end{aligned}$$

■

Example to illustrate potential multiplicity of optimal promotion standard:

We provide an example to illustrate that there may be multiple solutions to the condition

$$(1+S)\left(c_2 - c_1 + (d_2 - d_1)\, a_l^p\right) = w_{l2}(2) - w_{l2}(1),$$

which can be restated as

$$\begin{aligned} & (1+S)(d_2 - d_1)\, a_l^p + S(c_2 - c_1) \\ = \; & \frac{d_2}{1 - F_l(a_l^p)} \int_{a_l^p}^{\bar{a}_l} a_{jl} f_l(a_{jl})\, da_{jl} - \frac{d_1}{F_l(a_l^p)} \int_{\underline{a}_l}^{a_l^p} a_{jl} f_l(a_{jl})\, da_{jl}. \end{aligned}$$

In the example, we make the following assumptions regarding the parameters:

$$\underline{a}_l = c_2 = 0, d_1 = 0.00001, \bar{a}_l = c_1 = 1, d_2 = 1.57, S = S_1 = 0.57.$$

We assume a piecewise uniform distribution for a_{jl}, given by

$$f_l(a_{jl}) = \begin{cases} 4.99993 \text{ if } a_{jl} \in [0, 0.1], \\ 2.5 \text{ if } a_{jl} \in [0.7, 0.85] \cup [0.95, 1], \\ 0.00001, \text{ otherwise,} \end{cases}$$

implying

$$F_l(a_{jl}) = \begin{cases} 4.99993 a_{jl} \text{ if } a_{jl} \in [0, 0.1], \\ 0.00001 a_{jl} + 0.499992 \text{ if } a_{jl} \in [0.1, 0.7], \\ 2.5 a_{jl} - 1.250001 \text{ if } a_{jl} \in [0.7, 0.85], \\ 0.00001 a_{jl} + 0.8749905 \text{ if } a_{jl} \in [0.85, 0.95], \\ 2.5 a_{jl} - 1.5 \text{ if } a_{jl} \in [0.95, 1]. \end{cases}$$

The above condition then simplifies to

$$1.57 \cdot 1.56999 a_l^p - 0.57 - \frac{1.57}{1 - F_l(a_l^p)} \int_{a_l^p}^{1} a_{jl} f_l(a_{jl})\, da_{jl} + \frac{0.00001}{F_l(a_l^p)} \int_{0}^{a_l^p} a_{jl} f_l(a_{jl})\, da_{jl} = 0.$$

If we assume $a_l^p \in [0.7, 0.85]$, the condition becomes

$$\begin{aligned} & 1.57 \cdot 1.56999 \cdot a_l^p - 0.57 \\ & - \frac{1.57}{2.250001 - 2.5 a_l^p} \left(\int_{a_l^p}^{0.85} 2.5 a_{jl} da_{jl} + \int_{0.85}^{0.95} 0.00001 a_{jl} da_{jl} + \int_{0.95}^{1} 2.5 a_{jl} da_{jl} \right) \\ & + \frac{0.00001}{2.5 a_l^p - 1.250001} \left(\int_{0}^{0.1} 4.99993 a_{jl} da_{jl} + \int_{0.1}^{0.7} 0.00001 a_{jl} da_{jl} + \int_{0.7}^{a_l^p} 2.5 a_{jl} da_{jl} \right) = 0, \end{aligned}$$

which has the two solutions $a_{l,1}^p = 0.84531$ and $a_{l,2}^p = 0.81456$.

References

Aigner, Dennis J., and Glen G. Cain. 1977. Statistical theories of discrimination in labor markets. *Industrial and Labor Relations Review* 30, 175-187.

Al-Najjar, Nabil I. 2004. Aggregation and the law of large numbers in large economies. *Games and Economic Behavior* 47, 1-35.

Andre, Claire, Velasquez, Manuel, and Tim Mazur. 1992. Affirmative Action: Twenty-five Years of Controversy. *Issues in Ethics* 5.

Athey, Susan, Avery, Christopher, and Peter Zemsky. 2000. Mentoring and diversity. *American Economic Review* 90: 765-786.

Balafoutas, Lukas, and Matthias Sutter. 2012. Affirmative action policies promote women and do not harm efficiency in the laboratory. *Science* 335: 579-582.

Becker, Gary S. 1957. The economics of discrimination. University of Chicago Press, Chicago.

Bernhardt, Dan. 1995. Strategic promotion and compensation. *Review of Economic Studies* 62: 315-339.

Bertrand, Marianne, Black, Sandra E., Jensen, Sissel, and Adriana Lleras-Muney. 2014. Breaking the glass ceiling? The effect of board quotas on female labor market outcomes in Norway. *NBER Working Paper* No. 20256.

Bjerk, David. 2008. Glass ceilings or sticky floors? Statistical discrimination in a dynamic model of hiring and promotion. *The Economic Journal* 118: 961-982.

Black, Dan A.. 1995. Discrimination in an equilibrium search model. *Journal of Labor Economics* 13: 309-333.

Bognanno, Michael L., and Eduardo Melero. 2016. Promotion signals, experience and education. *Journal of Economics and Management Strategy* 25: 111-132.

Brown, Alasdair, and Subhasish M. Chowdhury. 2017. The hidden perils of affirmative action: sabotage in handicap contests. *Journal of Economic Behavior & Organization* 133: 273-284.

Burger, Rulof, and Rachel Jafta. 2012. Affirmative action in South Africa: an empirical assessment of the impact on labour market outcomes. In *Affirmative action in plural societies*, edited by Graham Brown, Arnim Langer and Frances Stewart, 80-99, Basingstoke, UK: Palgrave.

Calsamiglia, Caterina, Franke, Jörg, and Pedro Rey-Biel. 2013. The incentive effects of affirmative action in a real-effort tournament. *Journal of Public Economics* 98: 15-31.

Cassidy, Hugh, DeVaro, Jed, and Antti Kauhanen. 2016. Promotion signaling, gender, and turnover: New theory and evidence. *Journal of Economic Behavior & Organization* 126: 140-166.

Coate, Stephen, and Glenn C. Loury. 1993a. Antidiscrimination enforcement and the problem of patronization. *American Economic Review* 83: 92-98.

Coate, Stephen, and Glenn C. Loury. 1993b. Will affirmative-action policies eliminate negative stereotypes?. *American Economic Review* 83: 1220-1240.

Dato, Simon, Grunewald, Andreas, Kräkel, Matthias, and Daniel Müller. 2016. Asymmetric employer information, promotions, and the wage policy of firms. *Games and Economic Behavior* 100, 273-300.

DeVaro, Jed, Ghosh, Suman, and Cindy Zoghi. 2012. Job characteristics and labour market discrimination in promotions: new theory and empirical evidence. Discussion paper.

DeVaro, Jed, and Antti Kauhanen. 2016. An "Opposing Responses" Test of Classic versus Market-Based Promotion Tournaments. *Journal of Labor Economics* 34: 747-779.

DeVaro, Jed, and Michael Waldman. 2012. The signaling role of promotions: further theory and empirical evidence. *Journal of Labor Economics* 30: 91-147.

Ewerhart, Christian. 2014. Unique equilibrium in rent-seeking contests with a continuum of types. *Economics Letters* 125: 115-118.

Estevan, Fernanda, Gall, Thomas, Legros, Patrick, and Andrew F. Newman. 2014. College admission and high school integration. *Department of Economics - FEA/USP Working paper* No. 2014-26.

Fang, Hanming, and Andrea Moro. 2011. Theories of statistical discrimination and affirmative action: a survey. In *Handbook of Social Economics*, IA, edited by Jess Benhabib, Alberto Bisin, and Matthew Jackson, 133-200.

Fang, Hanming, and Peter Norman. 2006. Government-mandated discriminatory policies: theory and evidence. *International Economic Review* 47: 361-389.

Fryer, Roland G. 2007. Belief flipping in a dynamic model of statistical discrimination. *Journal of Public Economics* 91: 1151-1166.

Fryer, Roland G., and Glenn C. Loury. 2005. Affirmative action and its mythology. *Journal of Economic Perspectives* 19: 147-162.

Fu, Qiang. 2006. A theory of affirmative action in college admissions. *Economic Inquiry* 44: 420-428.

Ghosh, Suman, and Michael Waldman. 2010. Standard promotion practices versus up-or-out contracts. *RAND Journal of Economics* 41: 301-325.

Greenwald, Bruce C. 1986. Adverse selection in the labor market. *Review of Economic Studies* 53: 325-347.

Gürtler, Marc, and Oliver Gürtler. 2015. The optimality of heterogeneous tournaments. *Journal of Labor Economics* 33: 1007-1042.

Holmström, Bengt. 1982. Managerial incentive problems – a dynamic perspective. In *Essays in Economics and Management in Honor of Lars Wahlbeck*. Helsinki. Reprinted in *Review of Economic Studies* 66: January 1999, 169-82.

Holzer, Harry, and David Neumark. 2000. Assessing affirmative action. *Journal of Economic Literature* 38: 483-568.

Konrad, Kai A. 2009. Strategy and dynamics in contests. Oxford University Press, Oxford.

Krishnamurthy, Prasad, and Aaron Edlin. 2014. Affirmative action and stereotypes in higher education admissions. *NBER Working Paper* No. 20629.

Kurtulus, Fidan Ana. 2012. Affirmative action and the occupational advancement of minorities and women during 1973-2003. *Industrial Relations* 51: 213-246.

Lang, Kevin, and Jee-Yeon K. Lehmann. 2012. Racial discrimination in the labor market: theory and empirics. *Journal of Economic Literature* 50: 959-1006.

Lazear, Edward P., and Sherwin Rosen. 1981. Rank-order tournaments as optimum labor contracts. *Journal of Political Economy* 89: 841-864.

Leonard, Jonathan S. 1996. Wage disparities and affirmative action in the 1980's. *American Economic Review* 86: 285-289.

Milgrom, Paul and Sharon Oster. 1987. Job discrimination, market forces, and the invisibility hypothesis. *Quarterly Journal of Economics* 102: 453-476.

Moro, Andrea, and Peter Norman. 2003. Affirmative action in a competitive economy. *Journal of Public Economics* 87: 567-594.

Owan, Hideo. 2004. Promotion, turnover, earnings, and firm-sponsored training. *Journal of Labor Economics* 22: 955-978.

Phelps, Edmund S. 1972. The statistical theory of racism and sexism. *American Economic Review* 62: 659-661.

Pinkston, Joshua C. 2009. A model of asymmetric employer learning with testable implications. *Review of Economic Studies* 76: 367-394.

Schönberg, Uta. 2007. Testing for asymmetric employer learning. *Journal of Labor Economics* 25: 651-691.

Schotter, Andrew, and Keith Weigelt. 1992. Asymmetric tournaments, equal opportunity laws, and affirmative ction: some experimental results. *Quarterly Journal of Economics* 107: 511-539.

Shankar, Kameshwari. 2016. Multi-skilling, specialization and lateral job mobility. Discussion paper.

Smith, James P., and Finis Welch. 1984. Affirmative action and labor markets. *Journal of Labor Economics* 2: 269-301.

Waldman, Michael. 1984. Job assignments, signalling, and efficiency. *RAND Journal of Economics* 15: 255-267.

Waldman, Michael. 2013. Classic promotion tournaments versus market-based tournaments. *International Journal of Industrial Organization* 31: 198-210.

Waldman, Michael. 2016. The dual avenues of labor market signaling. *Labour Economics* 41: 120-134.

Waldman, Michael, and Ori Zax. 2016. An exploration of the promotion signaling distortion. *Journal of Law, Economics & Organization* 32: 119-149.

Zábojník, Ján, and Dan Bernhardt. 2001. Corporate tournaments, human capital acquisition, and the firm size-wage relation. *Review of Economic Studies* 68: 693-716.

Zábojník, Ján. 2012. Promotion tournaments in market equilibrium. *Economic Theory* 51: 213-240.