On Why Affirmative Action May Never End and Why it Should

Philippe Jehiel, Matthew Leduc

PARIS SCHOOL OF ECONOMICS
ECOLE D'ECONOMIE DE PARIS

WORKING PAPER N° 2021 – 54

On Why Affirmative Action May Never End and Why it Should

Philippe Jehiel
Matthew V. Leduc

JEL Codes: D40; I28; I30; J15
Keywords: Affirmative Action, General Equilibrium, Loss Aversion, Prospect Theory, Moral Hazard, Game Theory.

Funded by a French government subsidy managed by the ANR under the framework of the Investissements d'avenir programme reference ANR-17-EURE-001

On Why Affirmative Action May Never End and Why it Should*

Philippe Jehiel†
Paris School of Economics
& University College London

Matthew V. Leduc‡
Paris School of Economics
& Université Paris 1

August 9, 2021

Abstract

Successive governments must decide whether to implement an affirmative action policy aimed at improving the performance distribution of the next generation of a targeted group. Workers receive wages corresponding to their expected performance, suffer a feeling of injustice when getting less than their performance, and employers do not (perfectly) observe whether workers benefited from affirmative action. We find that welfare-maximizing governments choose to implement affirmative action *perpetually*, despite the resulting feeling of injustice that eventually dominates the purported beneficial effect on the performance of the targeted group. This is in contrast with the first-best that requires affirmative action to be temporary.

Keywords: Affirmative Action, General Equilibrium, Loss Aversion, Prospect Theory, Moral Hazard, Game Theory

JEL codes: D40; I28; I30; J15

*Jehiel thanks the ERC (grant no 742816) for funding. We thank seminar participants at Université Paris 1 and the Paris School of Economics, namely Agniezka Rusinowska, Emily Tanimura, Gabrielle Demange, Jean-Marc Tallon, Leonardo Pejsachowicz, Francis Bloch, Frédéric Koessler, Olivier Tercieux. We also thank Matthew L. Elliott and Larry Samuelson for useful comments.

†Email: philippe.jehiel@gmail.com

‡(Corresponding author) Email: mattvleduc@gmail.com. Mailing address: Paris Jourdan Sciences Economiques, 48 boulevard Jourdan, 75014 Paris, France. Tel.: (+33) 01 80 52 16 60

"I yield to no one in my earnest hope that the time will come when an affirmative action program is unnecessary and is, in truth, only a relic of the past. [...] within a decade at most, American society must and will reach a stage of maturity where acting along this line is no longer necessary."

Supreme Court justice Harry Blackmun, 1978

"The Court expects that 25 years from now, the use of racial preferences will no longer be necessary to further the interest approved today."

Supreme Court justice Sandra Day O'Connor, 2003

1 Introduction

The original rationale for affirmative action was to help underrepresented groups close achievement gaps and it was meant to be temporary. Decades after their inception, affirmative action policies however often remain in place. This article will attempt to provide an explanation by studying the incentives of successive governments to implement affirmative action policies.

The existing literature on affirmative action is vast and often tries to describe or explain inequalities between groups. Early developments include taste-based theories of discrimination (e.g. Becker (1957)), which suppose that exogenous preferences generate wage differences between groups, although the latter are unlikely to persist in competitive markets. Statistical discrimination theories, on the other hand, mainly attempt to explain outcome differences using imperfect information about the workers' performance levels, which leads to different wages being rationally paid to workers of different groups (e.g. Phelps (1972), Arrow (1973), Lundberg and Startz (1983), Coate and Loury (1993a) or Coate and Loury (1993b)). Such models often also link these different wages to the workers' incentives to invest in human capital, thus sustaining a performance gap between groups.

We depart from this literature, as we are not concerned in the current article with the origins of performance gaps between different groups. Instead, we employ a reduced-form approach in which we take the performance gap between two groups to be initially given, irrespectively of its origin. On the other hand, we analyze a setting in which governments believe that an affirmative action policy improves the performance distribution of an underrepresented (targeted) group of workers in future periods. This belief, held by governments, is in line with popular role model theories (see, for example, Chung (2000)), according to which witnessing certain members of an underepresented group achieving success would lead other group members to achieve higher success in the future.

We thus adopt a welfarist approach and seek to study the incentives of successive governments to implement affirmative action policies during their tenure. We do so using a repeated game setting, with each successive government trying to maximize a weighted sum of the welfare of two groups: the main group – which can be thought of as a majority group, even if we make no specific assumption on this group being larger in our formal model – and a targeted group – which can be thought of as a minority group. The main group, called A, initially has a distribution of performance that is different from that of the targeted group, called B. In applications, we will have in mind that the distribution for group A initially first-order stochastically dominates that of group B, even though this plays no role in our conclusions. As said earlier, this is a reduced-form approach and we do not try to explain the initial performance gap.

In the main part of the paper, we use a stylized model to express our arguments as clearly as possible. We later explain how its assumptions can be mostly relaxed without affecting our main conclusions. We suppose that employers cannot condition wages on group membership, which is in line with many anti-discrimination policies. In a perfectly competitive labor market, each employer pays a worker a wage equal to his expected performance, but does not observe whether the worker benefited from affirmative action or not and can only estimate this performance based on a curriculum vitae, which may be artificially improved by affirmative action. Paying workers a wage equal to their expected performance thus means that non-beneficiaries of affirmative action (of either the main or the targeted group) will get a wage below their true performance level. We postulate that in such a case, the worker suffers from a *feeling of injustice* that is proportional to the difference between his true performance (which the worker knows) and his wage. We believe such a feeling of injustice is very common, as suggested by recent so-called "populist" reactions. Note that this depressed wage can be understood, in a broader context, as being associated with the devaluation of a worker's diplomas or career promotions, which results from the possibility that he may have benefited from affirmative action.

In a first-best scenario, this depressed wage given to non-beneficiaries of affirmative action (and the associated feeling of injustice) means that affirmative action should not last permanently. The optimal duration would be determined, namely, by the weights governments place on the welfare of the main and the targeted groups, respectively. Indeed, it is essentially a tradeoff between the purported benefits to the targeted group, which stem from its improved performance distribution, and the feeling of injustice suffered by the non-beneficiaries (of either the main or the targeted group), which stems from the depressed wage. If governments place relatively more weight on the welfare of the main group, then affirmative action would end more quickly. It they place relatively more weight on the welfare of the targeted group, then it would last longer. However, as long as the governments care no less about the welfare of the main group than that of the targeted group, affirmative action would necessarily be ended at some point, as soon as the main group suffers some (even very small) feeling of injustice.

To the contrary, we show that successive governments *always* choosing to implement an affirmative action policy is an equilibrium (and is the unique equilibrium under some extra conditions). The intuition is that whether a government actually implements an affirmative action policy is not observable (or is only imperfectly observable) by employers. Therefore a deviation towards implementing an affirmative action policy is perceived by a government as having no effect on decreasing wages, while it is also believed to improve the performance distribution of the targeted group (through a role model argument). This creates a moral hazard, by which each government necessarily chooses to implement an affirmative action policy and fails to internalize the effect that it has on devaluing diplomas and promotions (and thus on depressing wages). The fact that actual affirmative action policy decisions are not observed (or are only imperfectly observed) is justified by the fact that it is often very difficult in practice to determine whether a government actually implemented an affirmative action policy or not. For example, in the United States, these policies are complex, they vary from state to state and even when they are not officially implemented, they may actually take place through private channels (e.g. non-governmental diversity enhancement programs, etc.). Thus, an actual deviation from the official (equilibrium) policy may not be (perfectly) observed by employers in the labor market.

The paper is organized as follows. In Section 2, we introduce the basic setting and define the workers' utilities and welfare. In Section 3, we then study how employers set the wages they pay to workers and show that it leads to a feeling of injustice felt by non-beneficiaries of affirmative action. In Section 4, we analyze each government's welfare maximization problem and present the two central results: (i) perpetual affirmative action as an equilibrium policy and (ii) the first-best policy where affirmative action is ultimately ended. In Section 5, we discuss how our assumptions can be relaxed, as well as model extensions. We also compare our model with the existing literature. We namely discuss why the feeling of injustice, present in our paper, cannot be captured by existing simple models of statistical discrimination. All proofs, as well as some generalizations, are relegated to an appendix in Section 6.

2 Setting

A worker has a performance level $c \in [0,1]$. This can be understood, for instance, as his result to a standardized university admission test.

A population of workers consists of two groups: group A (the main group) and group B (the targeted group). Group A's performance distribution is initially assumed to differ[1] from that of group B. We also make no assumptions as to the causes of these different performance distributions. A whole literature has already focused on this matter (cf. Section 1). In our context, these different performance distributions can be seen as a reduced-form representation, which we use for other purposes that will be made clear in what follows.

At each time t, a different government must decide whether to implement an affirmative action policy for the duration of its tenure (one period). That is, it chooses an action $\sigma_t \in \{0,1\}$ (no affirmative action or affirmative action).

At any time t, group A's performance density is $f_A(c)$ while group B's performance density is $f_{B,n_t}(c)$, where $n_t = \sum_{s<t} \sigma_s$ is the number of times previous governments have implemented affirmative action policies. $f_A(c)$ and $f_{B,n_t}(c)$ have support $[0,1]$. We will describe later how $f_{B,n_t}(c)$ varies with n_t but intuitively as n_t increases, $f_{B,n_t}(c)$ shifts lower values of c to higher values, resulting in first-order stochastic dominance. Each agent lives for only one period. At each time t, with probability $\delta \in (0,1)$, a mass $|A|$ and a mass $|B|$ of new agents from groups A and B respectively are born, with performance levels drawn according to $f_A(c)$ and $f_{B,n_t}(c)$. We therefore call δ the population's survival probability for one period.[2]

2.1 Effect of affirmative action policy

An affirmative action policy has two effects. First, it gives an immediate artificial boost to the curriculum vitae of a worker benefitting from it. This models the fact that a beneficiary of affirmative action has expanded opportunities in terms of university admissions or career promotions compared to a non-beneficiary, thereby artificially enhancing the quality of his curriculum vitae. Second, it is also believed to have long-term, positive effects on the performance distribution of group B. This purported long-term effect is in line with popular "role model" theories (e.g. Chung (2000)). This second effect will be captured by the dependence of $f_{B,n_t}(c)$ on n_t.

[1]In the applications we will have in mind, it is reasonable to think that group A's performance distribution initially first-order stochastically dominates that of group B, although this will play no role in our analysis.

[2]The model can easily be extended to allow agents to live for more than one period and to have overlapping generations. Since such elaborations would play no role in our analysis, we have chosen the simpler setting in which agents just live for one period.

It is important to note that an affirmative action policy can be interpreted[3] as anything that artificially increases the quality of a curriculum vitae (immediate effect) and improves the performance distribution of future generations (purported role model effect).

2.1.1 Effect of affirmative action policy on curriculum vitae quality

When $\sigma_t = 1$, a member of group B with performance level $c \in [0,1]$, but who obtained his university degree under that government, will have a curriculum vitae quality $\overline{c} = g(c)$, where g is an increasing function such that $g(c) > c$, $\forall c \in (0,1)$, and $g(0) = 0$, $g(1) = 1$. The support of $\overline{c}$ is thus also $[0,1]$. An affirmative action policy therefore increases the curriculum vitae quality of a beneficiary above his actual performance level. By contrast, an affirmative action policy has no effect on the curriculum vitae quality of members of group A, nor on those of members of group B who did not benefit from the affirmative action policy (i.e. those who obtained their degree under a government that chose $\sigma_t = 0$). That is, their curriculum vitae quality corresponds to their actual performance level: $\overline{c} = c$.

2.1.2 Effect of affirmative action policy on actual performance

We suppose that if $\sigma_t = 1$, then the next period's performance distribution of group B is shifted so that $f_{B,n_{t+1}}(c) \succ f_{B,n_t}(c)$, where $\succ$ indicates strong first-order stochastic dominance[4]. Note that the effect of the shift is permanent, i.e. the improvement remains in all future periods. This purported improvement in the performance of future cohorts of workers is consistent with the "role model" argument.

If $\sigma_t = 1$ for all t, then $f_{B,n_t}(c) \uparrow \overline{f}_B(c)$. Since $f_{B,n_t}(c)$ converges from below to a limiting distribution $\overline{f}_B(c)$, this implies that the distributional improvements become smaller and smaller as governments keep implementing affirmative action policies. Group A's performance distribution $f_A(c)$ does not vary with t.

2.2 Utilities and welfare

A worker is of type $\theta = (c, \overline{c}, G)$, where c is his true performance level, $\overline{c}$ is his curriculum vitae quality and $G \in \{A, B\}$ is the group this worker belongs to. A worker knows his type and the wage function $\omega(\overline{c})$ set by employers, which is the wage the worker earns based on his curriculum vitae quality. This is formalized in the following definition.

Definition 1 *A wage function* $\omega : [0,1] \rightarrow [0,1]$ *determines the wage a worker earns when presenting a curriculum vitae of quality* $\overline{c}$ *to the employer.*

2.2.1 Utility

The utility of a type $(c, \overline{c}, G)$ worker is

$$u_G(\overline{c}, c) \quad = \quad \omega(\overline{c}) - \gamma_G \max\{c - \omega(\overline{c}), 0\} \tag{1}$$

[3]See Section 5.1.6 for a discussion of how our model can accommodate even more general interpretations of affirmative action.

[4]i.e. that $F_{B,n_t}(c) > F_{B,n_{t+1}}(c)$ for all $c \in (0,1)$.

where $\gamma_G \max\{c - \omega(\overline{c}), 0\}$, for some $\gamma_G > 0$, captures the fact that a feeling of "injustice" is suffered when a worker gets a salary that is below his true performance level. Note that we allow $\gamma_A \neq \gamma_B$ so as to capture that the feeling of injustice may differently affect groups A and B.

In particular, the utility of a type $(c, \overline{c}, A)$ worker is

$$u_A(\overline{c}, c) \quad = \quad \omega(c) - \gamma_A \max\{c - \omega(c), 0\}$$

since such a worker does not benefit from affirmative action and thus $\overline{c} = c$.

Similarly, the utility of a type $(c, \overline{c}, B)$ worker who did *not* benefit from affirmative action will be denoted as

$$u_B(\overline{c}, c|\{aa\}^{\mathsf{c}}) \quad = \quad \omega(c) - \gamma_B \max\{c - \omega(c), 0\}$$

since such a worker does not benefit from affirmative action and thus $\overline{c} = c$.

By contrast, the utility of a type $(c, \overline{c}, B)$ worker who benefited from affirmative action will be denoted

$$u_B(\overline{c}, c|\{aa\}) = \omega(\overline{c}) - \gamma_B \max\{c - \omega(\overline{c}), 0\}$$

and $c < \overline{c} = g(c)$ since such a worker does benefit from affirmative action.

In the above, we assume that workers have the correct perception of their performance level c. We also note that there are no extra positive effects on utility of receiving a wage greater than the performance level. Such an asymmetry in the utility assessment of wages above or below the performance level is in line with well documented psychological studies (see in particular the prospect theory of Kahneman and Tversky (1979)), which suggest a different assessment for payoff realizations above or below the reference point (here naturally identified with the performance level).

2.2.2 Welfare

The welfare of each group at time t is defined by taking the aggregate utility of that group. We thus have,

$$W_{A,t} = |A| \int_0^1 u_A(\overline{c}, c) f_A(c) dc$$

$$W_{B,t|\{aa\}^{\mathsf{c}}} = |B| \int_0^1 u_B(\overline{c}, c|\{aa\}^{\mathsf{c}}) f_{B,n_t}(c) dc$$

$$W_{B,t|\{aa\}} = |B| \int_0^1 u_B(\overline{c}, c|\{aa\}) f_{B,n_t}(c) dc$$

and thus

$$W_{B,t} = |B| \int_0^1 \Big(\sigma_t u_B(\overline{c}, c|\{aa\}) + (1 - \sigma_t) u_B(\overline{c}, c|\{aa\}^{\mathsf{c}})\Big) f_{B,n_t}(c) dc$$

where σ_t is the actual policy decision made by the time-t government.

3 Effect of affirmative action policy on wage levels

3.1 Informational environment

In the main part of the paper, we assume employers at time t do not observe the sequence $\{\sigma_s\}_{s=1}^{t}$ of actual policy decisions made by specific governments. Moreover, they do not observe the actual performance distribution $f_{B,n_t}(c)$, otherwise they would be able to infer the sum of the actual σ_t, i.e. $n_t = \sum_{s<t} \sigma_s$. For simplicity of exposition[5], we will also suppose that they do not observe the time t and thus do not know under which government the workers have obtained their qualifications. As a result, they cannot tell whether a worker benefited from affirmative action or not. In Sections 5.1.3 and 5.1.2, we note that our main insights still hold if employers observe the time t, as long as they do not *perfectly* observe σ_t.

As usual, in equilibrium employers know the strategy $\sigma^* = \{\sigma_s^*\}_{s=1}^{\infty}$ that is chosen by governments and thus they can compute the probability that a worker benefited from affirmative action, conditional upon observing his curriculum vitae quality $\bar{c}$. Thus employers are aware of the general policy on matters of affirmative action, without being able to actually observe (perfectly) whether it is implemented or not in a given period.

A justification for the unobservability (or imperfect observability) of the actual σ_t (which plays a key role in our results) is that it is often very difficult to observe whether a government actually implements an affirmative action policy. For example, in the United States, these policies vary from state to state and even when they are not officially implemented, they may actually take place through private channels (e.g. non-governmental diversity enhancement programs, etc.). Thus, an actual deviation from the official policy σ_t^* may not be observed by the employer.[6]

Likewise, the unobservability of the actual time t (which is not essential for our results) models the fact that it is very difficult for an employer to determine whether a worker benefited from affirmative action when he obtained his degree. This indeed depends on whether an affirmative action policy was in place the year the worker graduated from university. It also depends on where the worker obtained his degree (which university, state, etc.) and whether an affirmative action policy was implemented in that location. Keeping track of all this is difficult for an employer.

Finally, the unobservability of the performance distribution $f_{B,n_t}(c)$ can be justified by the fact that these are very difficult to estimate and studies published on that topic are often imprecise or time-lagged. The only thing that needs to be understood (or believed) by the employers is the purported[7] dependence of $f_{B,n_t}(c)$ on $n_t = \sum_{s<t} \sigma_s^*$.

3.2 Setting wages

We consider a perfectly competitive labor market, where an employer pays a worker a wage equal to his expected performance level. In Section 6.1 of the Appendix, this reduced-form approach is rationalized based on a Bertrand-type model of competition between employers.

We also assume, in the main part of the paper, that employers are not allowed to take group

[5]This assumption allows employers to set a time-independent wage function. In Section 5.1.3, we show that although lifting this assumption leads to a time-dependent wage function, it does not affect our results.

[6]As said earlier, we discuss later on in Section 5.1.2 the robustness of our insights if the labor market only imperfectly observes the choices of σ_t .

[7]Note that this dependence of $f_{B,n_t}(c)$ on $n_t = \sum_{s<t} \sigma_s^*$ could indeed simply be "believed". It is not necessary that it is actually taking place. Our results are robust to that.

information (A or B) into account when giving a wage to a particular worker (although we explain in Section 5.1.1 that this assumption can be relaxed under natural elaborations of the model). This is consistent with anti-discrimination laws enacted in many countries and occupational areas[8]. That is, they must set a wage conditioned only by the curriculum vitae quality $\bar{c}$. The wage $\omega^*(\bar{c})$ paid to a worker of type $(c, \bar{c}, A)$ or to a worker of type $(c, \bar{c}, B)$ is thus the conditional expectation $\mathbb{E}[c|\bar{c}, \sigma^*]$ of the worker's true performance level c, expressed in the following lemma.

Lemma 1 *Given an equilibrium government policy strategy σ^*, the wage paid to a worker with curriculum vitae quality $\bar{c}$ has the form*

$$\omega^*(\bar{c}) = \sum_{t=1}^{\infty} \mathbb{P}(t) \left(\mathbb{P}_t^*(\{aa\}|\bar{c}) \cdot g^{-1}(\bar{c}) + \left(1 - \mathbb{P}_t^*(\{aa\}|\bar{c})\right) \cdot \bar{c} \right)$$

where

$$\mathbb{P}_t^*(\{aa\}|\bar{c}) = \frac{|B|\sigma_t^* f_{B,n_t}(g^{-1}(\bar{c}))/g'^{-1}(\bar{c}))}{|A|f_A(\bar{c}) + |B|(1-\sigma_t^*)f_{B,n_t}(\bar{c}) + |B|\sigma_t^* f_{B,n_t}(g^{-1}(\bar{c}))/g'^{-1}(\bar{c}))}$$

and $\{aa\}$ is the event that a worker benefited from affirmative action while

$$\mathbb{P}(t) = \delta^{t-1}(1-\delta)$$

is the probability of being at time t.

In words, $\omega^*(\bar{c})$ is a convex combination between $g^{-1}(\bar{c})$ and $\bar{c}$, where the weight assigned to $g^{-1}(\bar{c})$ should be proportional to the overall probability that a worker with curriculum vitae $\bar{c}$ benefited from affirmative action. This overall probability can be decomposed into a probability of benefiting from affirmative action conditional on $\bar{c}$ at the various times t (taking into account the chosen policy σ_t^* at t, hence the expression for $\mathbb{P}_t^*(\{aa\}|\bar{c})$) and the probability of being at time t as reflected by $\mathbb{P}(t)$.

We make the following assumption on $\mathbb{E}[c|\bar{c}, \sigma^*]$ for simplicity of exposition. Our results do not depend on it, but it will allow us to present them in a simpler manner, since we can rule out strategic behavior by which an agent could present a curriculum vitae of lower[9] quality than $\bar{c}$. An extension where a worker is allowed to present a curriculum vitae of a different quality than $\bar{c}$ is presented in the Appendix, where the robustness of our results to such strategic behavior is established in a more general context.

Assumption 1 *The conditional expectation $\mathbb{E}[c|\bar{c}, \sigma^*]$, and thus the wage function $\omega^*(\bar{c})$, is non-decreasing in $\bar{c}$.*

This assumption is easily satisfied under some conditions, e.g. when the likelihood ratio $\frac{f_A(c)}{f_B(g(c))}$ is increasing or when the mass $|A|$ is sufficiently larger than the mass $|B|$.

Relying on the expression of equilibrium wage derived in Lemma 1, we make the simple observation that whether the earned wage lies above or below the performance level solely depends on whether or not the worker benefitted from affirmative action.

[8]We show in Section 5.1.1, that our results actually hold even if group-based discrimination is allowed, as long as some members of the targeted group B do not benefit from affirmative action (a fairly weak assumption). In such a case, the feeling of injustice is suffered entirely by them (and not by group A members) and this is enough for our results to hold.

[9]Indeed, if the wage function $\omega^*(\bar{c})$ is decreasing on some parts of the support $[0, 1]$, a worker could earn a higher wage by presenting a curriculum vitae of lower quality than $\bar{c}$.

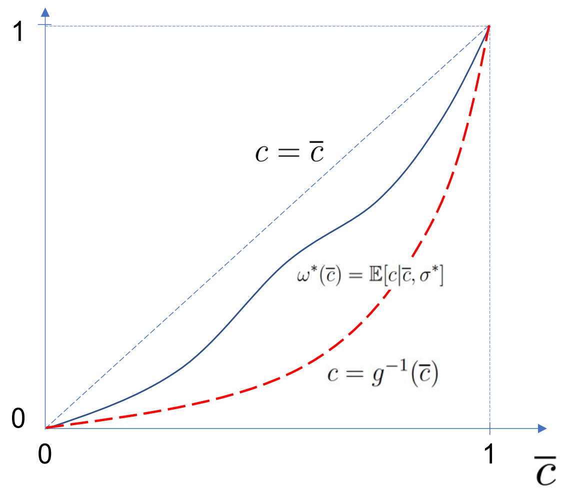

Figure 1: Illustration of Lemma 2: The wage $\omega^*(\overline{c})$ (full blue curve) is lower than the performance level of a non-beneficiary (thinly dotted blue curve) and higher than that of a beneficiary of affirmative action (thickly dotted red curve).

Lemma 2 *If a worker benefits from affirmative action (i.e. $c = g^{-1}(\overline{c})$), then he gets a wage higher than his performance level (i.e. $c < \omega^*(\overline{c})$). If a worker does not benefit from affirmative action (i.e. $c = \overline{c}$), then he gets a wage lower than his performance level (i.e. $c > \omega^*(\overline{c})$).*

Lemma 2 is illustrated in Fig. 1. Using Lemma 2, we now make the following observations.

Observation 1 *Given the equilibrium wage function $\omega^*(\overline{c})$, the utility of a member of group A can be written as*

$$\begin{aligned} u_A(\overline{c}, c) &= \omega^*(c) - \gamma_A \max\{c - \omega^*(c), 0\} \\ &= \omega^*(c) - \gamma_A(c - \omega^*(c)) \end{aligned}$$

since $\overline{c} = c$ and $\omega^*(c) < c$ for a member of group A (as he does not benefit from affirmative action). Such a worker suffers a feeling of injustice.

Observation 2 *Given the equilibrium wage function $\omega^*(\overline{c})$, the utility of a member of group B who does* not *benefit from affirmative action can be written as*

$$\begin{aligned} u_B(\overline{c}, c|\{aa\}^{\mathsf{c}}) &= \omega^*(c) - \gamma_B \max\{c - \omega^*(c), 0\} \\ &= \omega^*(c) - \gamma_B(c - \omega^*(c)) \end{aligned}$$

since $\overline{c} = c$ and $\omega^*(c) < c$ for a member of group B not benefiting from affirmative action. Such a worker also suffers a feeling of injustice.

Observation 3 *Given the equilibrium wage function $\omega^*(\overline{c})$, the utility of a member of group B who benefits from affirmative action can be written as*

$$\begin{aligned} u_B(\overline{c}, c|\{aa\}) &= \omega^*(\overline{c}) - \gamma_B \max\{c - \omega^*(\overline{c}), 0\} \\ &= \omega^*(\overline{c}) \end{aligned}$$

since $\overline{c} = g(c) > c$ and $\omega^*(\overline{c}) > c$ for a member of group B benefiting from affirmative action. Such a worker does not suffer a feeling of injustice.

3.2.1 A broader interpretation of the depressed wage

In our model, the wage is depressed due to the possibility that a worker benefited from affirmative action. This represents the fact that a certain curriculum vitae quality is, in expectation, no longer associated to the same performance level as if there were no affirmative action policy. Indeed, an affirmative action policy has the effect of devaluing the diplomas or promotions that figure on a worker's curriculum vitae, if there is only some chance that the worker may have benefited from such a policy.

4 The government's decision problem

4.1 Informational environment

For simplicity of exposition, we will suppose that a time-t government knows the time t. It is however irrelevant[10] whether it knows t and our results do not depend on this. It is also inessential whether a time-t government observes the past policies $\sigma_{t'}$, $t' < t$, that were actually chosen by other governments or the actual performance distribution $f_{B,n_t}(c)$ prevailing at time t. This will have no impact on its decision to implement an affirmative action policy or not, as it will become clear later. The government, however knows (or believes) the purported effect of n_t on $f_{B,n_t}(c)$ and thus believes that choosing an affirmative action policy $\sigma_t = 1$ will improve the performance distribution: $f_{B,n_{t+1}}(c) \succ f_{B,n_t}(c)$.

4.2 Policy decisions

For all $t \geq 1$, a time-t government wants to maximize the following objective function:

$$\max_{\sigma_t \in \{0,1\}} \quad \textstyle\sum_{s=t}^{\infty} (W_{A,s} + \lambda_B W_{B,s}) \delta^{s-t} \tag{2}$$

Here $\lambda_B \geq 0$ is a weight placed on the welfare of group B (the weight placed on the welfare of group A is normalized to 1). We suppose this weight is constant through time. $\delta \in (0, 1)$ is the population survival rate previously introduced and δ serves also as discount factor for the overall welfare determination.[11]

We will mostly consider the case when $\lambda_B \leq 1$. In particular, $\lambda_B = 1$ corresponds to the standard total welfare criterion and $\lambda_B < 1$ reflects a preference for the main group A in the governmental objective. We will also comment on the case when $\lambda_B > 1$, which reflects a preference for the targeted group B.

Since a time-t government knows the equilibrium $\sigma^* = \{\sigma_s^*\}_{s=1}^{\infty}$ that is played by other governments, it is able to compute $W_{A,s}$ and $W_{B,s}$, for $s > t$. In other words, $\omega^*(\overline{c})$ and $f_{B,n_s}(c)$ are taken to be consistent with this equilibrium play.

[10] This is further discussed in Section 5.1.4.

[11] One might consider additional factors (such as impatience) affecting the discount rate. To the extent that the employers and the governments use the same discount rate, our analysis and main insights remain unaffected.

Note that this welfare maximization problem can be interpreted either (i) as the governments caring about the workers or (ii) caring simultaneously about the workers and the employers. Indeed, the employers' welfare is their aggregate profit, which is 0 since we study a perfectly competitive Bertrand-style labor market where employers pay workers a wage equal to their expected productivity (i.e. expected performance level). The reader is referred to Section 6.1 of the Appendix for a more precise formulation of this labor market with Bertrand competition.

Our first main result, Proposition 1, states that different governments choosing to *perpetually* implement an affirmative action policy is always an equilibrium. We also give sufficient (but non necessary) conditions under which it is the unique equilibrium.

Proposition 1 (Equilibrium policy) *Given any $\lambda_B > 0$:*

(i) $\sigma_t^* = 1$ for all t is an equilibrium.

(ii) There exists $\overline{\gamma}_B$ such that for any $\gamma_B < \overline{\gamma}_B$, it is the unique equilibrium[12].

(iii) There exists $\overline{\beta}$ such that for any $\frac{|B|}{|A|+|B|} < \overline{\beta}$, it is the unique equilibrium.

The intuition behind Proposition 1(i) is that any government believes that implementing an affirmative action policy improves the performance distribution of future cohorts of workers. Moreover, since a policy deviation is not observed by employers, it cannot have any improving impact on the wage function ω^* chosen by employers. Therefore, there is no reason why a particular government would deviate from an equilibrium $\sigma_t^* = 1$ in which it implements an affirmative action policy.

The intuition behind Proposition 1(ii) is that a deviation from a putative equilibrium in which $\sigma_t^* = 0$ to $\sigma_t = 1$ would increase the average performance of future cohorts of B workers (and thus the average wage they receive), but could also potentially increase the average feeling of injustice felt by B workers not benefiting from affirmative action in future periods. Indeed, the feeling of injustice could worsen following an increase in the performance level, if the latter increases faster than the wage received at a higher performance level (recall that the feeling of injustice is $\gamma_B \max\{c - \omega^*(\overline{c}), 0\}$). A sufficient condition for the postive effect to dominate the negative one is that the parameter γ_B be small enough.

Finally, the intuition behind Proposition 1(iii) is very simple. As the fraction of group B workers in the population decreases, the probability that a worker is a beneficiary of affirmative action also decreases. The conditional expectation of a worker's performance $\mathbb{E}[c|\overline{c}, \sigma^*]$ (and thus the wage function $\omega^*(\overline{c})$) therefore converges to the actual worker's performance c, and the feeling of injustice of group B workers, $\gamma_B \max\{c - \omega^*(\overline{c}), 0\}$, necessarily decreases. By the same argument as in Part (ii), this is thus a sufficient condition for any deviation from a putative equilibrium where $\sigma_t^* = 0$ to $\sigma_t = 1$ to improve the welfare.

Our second main result, Proposition 2, states that in the first-best scenario, affirmative action policies *always* end after a finite number of periods.

Proposition 2 (First-best policy) *Suppose that at time $t = 0$, a single government announces (and commits to) the policy plan $\hat{\sigma} = \{\hat{\sigma}_t\}_{t=1}^{\infty}$ that maximizes the welfare function $\sum_{t=1}^{\infty} \delta^t (W_{A,t} +$*

[12] It is interesting to note that it is enough that such a parameter γ_B, capturing the feeling of injustice felt by members of group B not benefiting from affirmative action, corresponds to one chosen by the government and it need not be the actual one felt in population B. Indeed, recall that the equilibrium wage ω^* actually does not depend on γ_B. Only the welfare $W_{B,s}$ of group B does.

$\lambda_B W_{B,t})$*, and assume* $\gamma_A \neq 0$*. Then for any* $\lambda_B \in [0,1]$*, there exists* $\overline{\delta} \in (0,1)$ *such that for all* $\delta \in (\overline{\delta}, 1)$*,* $\{\hat{\sigma}_t\}_{t=1}^{\infty}$ *has a threshold form* $\hat{\sigma}_t = 1$ *for* $t < \overline{T}$ *and* $\hat{\sigma}_t = 0$ *for* $t \geq \overline{T}$*, for some (finite)* $\overline{T} \in \mathbb{N}$*.*

Proposition 2 essentially means that if different governments were able to coordinate their actions over time, they would never choose to make affirmative action permanent. The intuition is quite simple: After a certain number of periods the improvement in the performance distribution becomes marginal, while the depressing effect on wages is not. As a matter of fact, $f_{B,n_t}(c)$ converges from below to a limiting distribution $\overline{f}_B(c)$, implying that the distributional improvements become smaller and smaller as affirmative action policies are implemented over time.

The optimal threshold $\overline{T}$, while always finite, depends namely on the relative weight placed by governments on the welfare of the targeted group B relative to the main group A, i.e. on λ_B.

Note that when $\lambda_B < 1$ (i.e. when the government cares relatively more about group A than group B), the parameter γ_A governing the feeling of injustice of group A can be 0 and the first-best policy will still prescribe stopping affirmative action after a finite number of periods, because the depressed wage penalizes group A sufficiently while the performance distribution of group B is only marginally improved.

When $\lambda_B = 1$ (i.e. when the government cares equally about group A and group B), then since the average wage is equal to the average performance level across the market (i.e. $\mathbb{E}[\omega^*(\overline{c})] = \mathbb{E}[c]$), an affirmative action policy effectively represents just a transfer of welfare from group A to group B. Indeed, this transfer of welfare takes place through group A workers receiving wages lower than their performance levels while group B workers receive wages higher than their performance levels. In this case, as long as the parameter γ_A is *strictly greater* than 0 (no matter how small it is), a first-best policy will prescribe stopping affirmative action after a finite number of periods because otherwise the feeling of injustice felt by group A would become worse than the improvement in the performance distribution of group B after sufficiently many implementations of the affirmative action policy.

Finally, if λ_B were to be strictly greater than 1 (i.e. when the government cares relatively more about group B than group A), then we might need γ_A to be sufficiently positive in order to justify stopping affirmative action. In such a case, $\gamma_A \geq \underline{\gamma}_A$ for some $\underline{\gamma}_A > 0$ would be a sufficient (but not always necessary) condition for our first-best result to hold.

5 Some extensions

5.1 Discussion of assumptions

In the next subsections, we show that our main results are quite robust and most often hold, even if we relax the assumptions we have made in the main part of the paper. Namely, in Sections 5.1.1 to 5.1.4, we explain that all we really need for our main results to hold is that an employer cannot be certain that a group-B worker has not benefited from affirmative action.

In Section 5.1.5, we extend our model to include a labor market congestion externality, while in Section 5.1.6, we briefly discuss an alternative and complementary view about affirmative action in terms of a biased promotions process and suggest that it would give similar insights.

5.1.1 Allowing for group-based discrimination in the labor market

In our main model, we have not allowed the employers to condition the wages on the group A or B to which the worker belongs. This was motivated on the ground that such discrimination is in general forbidden. As should be clear, a conclusion similar to our main insights would hold if the amount of allowed discrimination was capped instead of being forbidden.[13] If on the other hand, discrimination were not regulated, then within our current model, governments would still choose $\sigma_t^* = 1$ in every period, but this time workers both from groups A and B would receive a wage equal to their performance levels. As a result, the equilibrium would coincide with the first-best.[14] Such a conclusion that unregulated discrimination leads to the first-best would however be fragile to the introduction of other elaborations. For example, if we assume that affirmative action would not reach the entire group B, then the share of group B that would not benefit from affirmative action would play a role similar to that of group A in the current model. While in equilibrium affirmative action would occur in every period, the first-best would require to stop affirmative action at some point so as to save on the feeling of injustice inflicted by affirmative action on the non-beneficiaries of group B. The general spirit of our finding is thus robust to richer specifications of the wage setting environment.

5.1.2 Making the policy σ_t partially observable to employers

From another perspective, in our main model, we have assumed that employers make no observation at all about the chosen policies σ_t. We note that if employers at t were to observe a noisy signal about σ_t (with a support of signal realization that would be the same whether $\sigma_t = 0$ or 1), then the same $\sigma_t^* = 1$ for all t would remain an equilibrium. This holds because in such an equilibrium, the signal observed by employers would be perceived to be uninformative and thus would have no effect on the chosen wage, thereby providing the incentive to governments to choose $\sigma_t = 1$ in all periods, as in the main model.[15] Since the first-best policy would be unaffected by this modification, this establishes the robustness of our main insights with respect to a broad class of observational environments.

5.1.3 Making the time period t observable to employers

In Section 3, we have made the assumption that the time t was unobservable to employers. This was argued to be justified in a real-life context, due to the difficulty for employers to determine whether a particular worker had benefited from affirmative action or not.

It is however important to note that this assumption is unnecessary to obtain our results. Indeed, if t were observable to employers, then they would also condition their expectation of a worker's

[13] One reason for banning (or regulating) discrimination not discussed in this paper is that it might trigger another feeling of injustice if the reference point of workers benefiting from affirmative action switches from the performance level c to the inflated level $\bar{c} = g(c)$ that results from affirmative action.

[14] It is worth noting that within our basic model, inefficiency arises because both affirmative action policies and anti-discrimination rules operate at the same time. This is in contrast to conventional views in which affirmative action and anti-discrimination rules are often regarded as necessary and complementary tools to achieve better equality and higher welfare.

[15] This observation is related to one made by Bagwell (1995) in the context of Stackelberg interactions in which the action of the first-mover would be observed with noise. He notes that the Nash equilibrium of the normal form game is then a Perfect Bayesian Equilibrium of the two-stage interaction. Van Damme and Hurkens (1997) later note that there are other equilibria involving mixed strategies, but such equilibria can be regarded as being less robust to the extent that they rely on indifferences of the players.

productivity on t and the wage function $\omega_t^*(\bar{c})$ would therefore depend on time.

Note that it is easy to find the expression of such a $\omega_t^*(\bar{c})$, simply by setting $\mathbb{P}(t) = 1$ and $\mathbb{P}(t') = 0$ for $t' \neq t$ in Lemma 1 (or, more generally, in Lemma 3 of the Appendix). Lemma 2 (or, more generally, Lemmas 4 and 5 of the Appendix) then holds for any period t in which affirmative action takes place ($\sigma_t^* = 1$) and thus a group A worker suffers a feeling of injustice in such a period. If affirmative action did not take place ($\sigma_t^* = 0$), then $\omega_t^*(\bar{c}) = c$ for all workers and, since all workers get a wage equal to their performance level, no feeling of injustice is suffered by anyone (neither in group A nor in group B) in such a period. Interestingly, this makes $\sigma_t^* = 1$ for all t the *unique* equilibrium for *any* γ_B and *any* $\frac{|B|}{|A|+|B|}$. Proposition 1 is then strengthened since the sufficient conditions for uniqueness given in Proposition 1(ii)-(iii) become irrelevant.

It is easy to check that Proposition 2 also still applies with a time-dependent salary $\omega_t^*(\bar{c})$ and the same intuition given in Section 4 explains the result.

Nevertheless, a model with a single, time-independent wage function $\omega^*(\bar{c})$ is simpler to present and the assumption of not observing t is also quite realistic.

5.1.4 Making the time period t unobservable to governments

Note that in Section 4, we supposed for simplicity of exposition that the time-t government observed t. This allowed for a neater expression of the welfare function in Eq. (2). This assumption is however unnecessary. Indeed, if the government does not observe t, then the welfare function simply becomes the expected welfare

$$\sum_{t=1}^{\infty} \mathbb{P}(t) \sum_{s=t}^{\infty} (W_{A,s} + \lambda_B W_{B,s}) \delta^{s-t} \tag{3}$$

where $\mathbb{P}(t)$ is the probability of being at time t, given in Lemma 1.

Thus, given some strategy profile σ, every particular time-t government maximizes the same expected welfare expressed in Eq. (3), by choosing a best-responding strategy $\sigma_t' \in \{0, 1\}$, which will necessarily be homogenous across governments. It follows that an equilibrium strategy profile $\sigma^* = \{\sigma_t^*\}_{t=1}^{\infty}$ is homogenous (i.e. all governments choose the same strategy, $\sigma_t^* = \sigma_{t'}^* \in \{0, 1\}$ for any $t \neq t'$).

Proposition 1 still holds. Indeed, since under the conditions of that proposition, it was the right decision for any time-t government to choose $\sigma_t = 1$, irrespectively of the actions of other governments (remember that a deviation to $\sigma_t = 1$ is unobserved by the labor market and thus has no effect on depressing wages while improving the performance distribution of group B), this remains true when a government does not know the time t.

Proposition 2 still applies without changes since it is a first-best result, where a single government chooses and commits to an optimal policy plan $\hat{\sigma} = \{\hat{\sigma}_t\}_{t=1}^{\infty}$ at $t = 0$.[16]

5.1.5 Adding a labor market congestion externality

In our model, non-beneficiaries of affirmative action suffer from receiving a wage that is lower than their actual performance level, while beneficiaries of affirmative action receive a wage that is higher

[16] If we forced σ_t to be time invariant, then the best commitment strategy would require randomization with a weight assigned to affirmative action that would tend to 0 as δ gets close to 1.

than their actual performance level. Therefore, a transfer of utility between groups arises through the wage channel.

From another perspective, affirmative action is often thought of as an allocation problem, e.g. allocating a finite number of jobs between two groups, which would result in extra transfers between beneficiaries and non-beneficiaries of affirmative action in addition to the wage effect considered in our main model. While modeling a full-scale matching process is beyond the scope of this paper, our model can be extended in such a direction by adding a labor market congestion externality. This will be represented by a positive term in the utility function of a beneficiary and a negative term in the utility function of a non-beneficiary.

Thus, the utility of an A worker will take the form

$$u_A(\bar{c}, c) = \omega^*(c) - \gamma_A(c - \omega^*(c)) - \frac{\eta}{|A|}\sigma_s$$

while that of a B worker benefiting from affirmative action will take the form

$$u_B(\bar{c}, c|\{aa\}) = \omega^*(\bar{c}) + \frac{\eta}{|B|}$$

where $\eta > 0$ is a parameter measuring the magnitude of the congestion externality. It is easy to see that the aggregate transfer of utility from group A to group B due to labor market congestion, in a period when an affirmative action policy is implemented (i.e. when $\sigma_s = 1$), is simply η.

A simple calculation shows that a time-t government's objective function (evaluated at the equilibrium σ^*) can now be written as

$$\sum_{s=t}^{\infty}(W_{A,s} + \lambda_B W_{B,s} + \sigma_s^*\eta(\lambda_B - 1))\delta^{s-t}$$

where $W_{A,s}$ and $W_{B,s}$ are the aggregate welfare of groups A and B at time s, *absent* the congestion externality (i.e. just as they were defined in Section 2.2.2), while the additional term $\sigma_s^*\eta(\lambda_B - 1)$ represents the welfare associated to the transfer of utility from group A to group B (in equilibrium) due to labor market congestion. This term will be negative if the weight λ_B assigned by the government to group B is smaller than 1.

We therefore have the following analogue of Proposition 1.

Observation 4 (Equilibrium policy with congestion)

(i) If $\lambda_B = 1$, then Proposition 1 follows[17].

(ii) If $\lambda_B \in [0, 1)$, there exists $\underline{\eta}$ such that $\forall \eta > \underline{\eta}$, then σ^* with a threshold form $\sigma_t^* = 1$ for $t < T_{eq}$ and $\sigma_t^* = 0$ for $t \geq T_{eq}$, for some (finite) $T_{eq} \in \mathbb{N}$ is an equilibrium.

The intuition behind part (i) is that, when $\lambda_B = 1$, then labor market congestion simply causes a welfare transfer between groups A and B since the government cares equally about the two groups and thus it has no impact on the aggregate welfare.

[17]Note that if λ_B were to be greater than 1, then a time-t government's decision to choose $\sigma_t = 1$ would be even further rewarded by the welfare transfer from group A to group B. Indeed, the only difference compared to Proposition 1 is that there would be an additional non-negative term $\sigma_t\eta(\lambda_B - 1)$ in the welfare function being maximized. Proposition 1 would thus also apply without change since whenever $\sigma_t = 1$ would have been chosen in the original setting, it would also be chosen in the presence of congestion.

The intuition behind part (ii) is that the benefits of choosing $\sigma_t = 1$ when sufficiently many governments have already done it before become marginal (due to the convergence of the performance distribution $f_{B,n_t}(c)$ to $\overline{f}_B(c)$). When $\lambda_B < 1$ and η is large enough, the negative term $\sigma_t\eta(\lambda_B - 1)$ in the welfare function then means that choosing $\sigma_t = 1$ will decrease the welfare more than it could increase due to a boosted beneficiary's curriculum vitae (i.e. $\bar{c} = g(c)$) and an improvement in $f_{B,n_t}(c)$. Naturally, the threshold T_{eq} is weakly decreasing in the magnitude of the congestion factor η.

We also have the following analogue of Proposition 2, in which we namely see that affirmative action will still last longer in equilibrium than a first-best policy would prescribe.

Observation 5 (First-best policy with congestion) *For any $\lambda_B \in [0, 1]$, Proposition 2 follows*[18]. *Moreover, the first-best threshold in the presence of congestion is weakly lower than the first-best threshold in the absence of congestion, i.e. $\overline{T}_{con} \leq \overline{T}$, and it is also weakly lower than the one of the equilibrium policy with congestion, i.e. $\overline{T}_{con} \leq T_{eq}$.*

The intuition behind the fact that the first-best threshold with congestion $\overline{T}_{con}$ is weakly lower than the first-best threshold without congestion $\overline{T}$ is that implementing an affirmative action policy ($\sigma_t = 1$) at any given time t has an additional penalty $\sigma_s^*\eta(\lambda_B - 1)$ when $\lambda_B < 1$. Thus, the benefits of improving the performance distribution of group B will be cancelled more quickly in the presence of congestion than in the case without congestion.

Similarly, the intuition behind that fact that the first-best threshold with congestion $\overline{T}_{con}$ is weakly lower than the equilibrium threshold with congestion T_{eq} is that in the first-best case, the single $t = 0$ government takes the depressing wage effects into account (in addition to the penalties $\sigma_s^*\eta(\lambda_B - 1)$) when choosing a multi-period affirmative action policy plan $\hat{\sigma}$. The threshold $\overline{T}_{con}$ at which affirmative action stops is thus necessarily weakly lower than in the equilibrium case.

5.1.6 Affirmative action as a biased promotion process

We have suggested throughout the paper that our basic model could accommodate different interpretations of the nature of affirmative action. In this subsection, we briefly elaborate further on one such interpretation. When affirmative action takes the form that a promotion at a given hierarchical level should favor someone from group B, this would translate to promoting someone of a potentially lower performance level to the extent that the pool of group B workers is typically smaller than the pool of group A workers. This would give rise to a value gap similar to the gap between $g^{-1}(c)$ and c discussed in the main model.[19] To the extent that anti-discrimination policies would force the wage associated to a given hierarchical level in an organization or firm to be independent of whether the employee comes from group A or B, the wage would be determined as an average between the value corresponding to an A employee and that corresponding to a B employee, giving rise to to a formula similar to that in the main model.

[18] If λ_B were to be strictly greater than 1 (i.e. when the government cares relatively more about group B than group A), then we might need the feeling of injustice parameter γ_A to be sufficiently positive in order to justify stopping affirmative action even earlier. In such a case, $\gamma_A \geq \underline{\gamma}_A$ for some $\underline{\gamma}_A > 0$ would be a sufficient (but not always necessary) condition for Observation 5 to hold.

[19] Considering the largest draws from finite samples of different sizes when coming from pool A or B would allow us to put more structure on the shape of the value gap (which our reduced form approach does not permit).

Finally, assuming the government has no direct control on the implementation of the affirmative action policy but can only affect it through non-transparent directives (as supported by the arguments discussed in the Introduction and throughout the paper), this would rationalize the fact that the actual affirmative action policy σ_t of the government in charge at time t would not be directly observable by the market, thereby giving the same incentives to governments to choose $\sigma_t = 1$ as in the main model. Thus, while such a view of affirmative action does not refer to a curriculum vitae boost but only to a biased promotion process, we should expect similar insights as in our main model: Affirmative action is implemented for too long as compared with the first-best.

5.2 Comparisons with existing literature

We mainly depart from the existing literature on affirmative action by studying the incentives of governments to implement affirmative action policies. Indeed, most of the literature focuses on other incentives: those linked to hiring decisions made by employers or to investments in human capital made by workers, which may be reduced by an affirmative action policy (e.g. Lundberg and Startz (1983), Coate and Loury (1993a) or Coate and Loury (1993b); see also Fang and Moro (2011) for a survey on discrimination and affirmative action).

Our argument for why governments have no incentive to stop implementing affirmative action policies is based on a new moral hazard argument: each government's actual policy decision is not observed by the employers in the labor market and thus has no effect on depressing wages. When choosing its policy, each government therefore has no incentive to internalize this depressing wage effect, which as we postulate in line with established psychological biases, creates a feeling of injustice among non-recipients of affirmative action. Naturally, since in equilibrium all successive governments choose to implement affirmative actions policies, then wages will be depressed and a feeling of injustice will be suffered by those who did not benefit from affirmative action. In practical applications, it is indeed reasonable to suppose that each government is likely to neglect the negative effect that an affirmative action policy could have on the employers' perception of a curriculum vitae in the labor market (and thus on wages).

Another key difference between our model and the statistical discrimination literature is that in the latter, wages are typically permitted to differ across groups. In such a case, the minority group receives a wage lower than that of the main (majority) group, because the minority group has coordinated with employers on an equilibrium with low investment in skills and low wage. In this literature, an affirmative action policy can allow this minority group to transit to the equilibrium with high investment in skills and high wage, without any adverse effects on the majority group. Thus, there is no equivalent to the feeling of injustice that appears in our model[20].

Moreover, many models in the existing literature also often suppose that performance distributions of both groups are initially the same and then try to explain how they could diverge in equilibrium. We need no such assumptions on initial distributions in our model, as they are not relevant to reach our conclusions. Our purpose is not in explaining the origins of outcome inequalities between groups, but rather the duration of affirmative action policies implemented by governments. The only distributional assumption required in our model is that the government believes that implementing an affirmative action policy will improve the performance distribution of the targeted

[20] For other models of affirmative action with endogenous wage determination, see for example Moro and Norman (2003).

group in future generations. Our setting thus lends itself well to popular "role model" arguments (e.g. Chung (2000)). These improved future performance distributions could result from positive peer effects within the group, which could for example encourage more investments in human capital, although none of this needs to be explicitly described as our model employs a reduced-form approach where performance distributions simply follow an improving dynamics.

6 Appendix

6.1 Wage setting with Bertrand competition

We suppose that each firm produces a numeraire good of price equal to 1 with a constant return to scale technology and using labor as the input. The quantity of the numeraire good produced by a unit mass of workers of performance level c is thus simply c. The profit generated by a unit mass of workers of performance level c, when they are paid a wage $\omega(\overline{c})$, is thus

$$\pi = c - \omega(\overline{c}).$$

Since a firm only observes the curriculum vitae quality $\overline{c}$ of a worker it hires, the expected profit generated by a unit mass of workers with such curriculum vitae is then

$$\mathbb{E}[\pi|\overline{c}] = \mathbb{E}[c|\overline{c}, \sigma^*] - \omega(\overline{c})$$

where, as we know, $\mathbb{E}[c|\overline{c}, \sigma^*]$ is the expected performance level of a worker presenting a curriculum vitae $\overline{c}$ when the equilibrium government affirmative action strategy is σ^*.

If the firm hires a mass q of workers with curriculum vitae qualities having a density function $f(\overline{c})$, then its expected profit is

$$\begin{aligned}
\Pi &= q\mathbb{E}[\pi] \\
&= q\int_{\overline{c}} \mathbb{E}[\pi|\overline{c}]f(\overline{c})d\overline{c} \\
&= q\int_{\overline{c}} \big(\mathbb{E}[c|\overline{c}, \sigma^*] - \omega(\overline{c})\big)f(\overline{c})d\overline{c} \qquad (4)
\end{aligned}$$

where Π is also the realized profit, since each worker has zero measure.

A firm will thus maximize this profit by choosing an optimal wage function ω^*. Note that the profit in Eq. (4) is additively separable across $\overline{c}$. A firm thus chooses, for each curriculum vitae quality $\overline{c}$, the wage $\omega^*(\overline{c})$ that maximizes

$$\mathbb{E}[\pi|\overline{c}] = \mathbb{E}[c|\overline{c}, \sigma^*] - \omega(\overline{c}).$$

Since we consider a perfectly competitive Bertrand setting, it follows that the optimal wage will be equal to a worker's expected performance level, i.e. $\omega^*(\overline{c}) = \mathbb{E}[c|\overline{c}, \sigma^*]$, which is the worker's marginal productivity. Indeed, giving a wage higher than $\mathbb{E}[c|\overline{c}, \sigma^*]$ would result in a negative profit from hiring workers of that curriculum vitae quality, while giving a wage lower than $\mathbb{E}[c|\overline{c}, \sigma^*]$ would result in another employer hiring the workers away with a slightly higher wage. It also follows that a firm's profit is zero, i.e. $\Pi = 0$, and thus even if the governments care about the firms' welfare, the latter will not appear in their objective function in Eq. (2).

6.2 A more general model

We present here a more general model, where agents can choose the curriculum vitae quality that they present to employers. This allows us to treat the more general case where the conditional expectation $\mathbb{E}[c|\hat{c}, \mu^*, \sigma^*]$ may not be monotone. We illustrate that the results presented previously still hold, since they are just a particular case of this more general setting.

In this general model, a wage function $\omega(\hat{c})$ set by employers is the wage the worker earns when declaring a curriculum vitae of quality $\hat{c} \in [0, 1]$ to the employer. Here, we see that a worker can declare a curriculum vitae of quality not necessarily equal to his actual quality $\overline{c}$. This is formalized in the following definition.

Definition 2 *A wage function $\omega : [0, 1] \to [0, 1]$ determines the wage a worker earns when declaring a curriculum vitae of quality $\hat{c}$ to the employer.*

The utility of a type $(c, \overline{c}, G)$ worker, when presenting a curriculum vitae of quality $\hat{c} \in [0, 1]$, is thus

$$u_G(\hat{c}, c) \quad = \quad \omega(\hat{c}) - \gamma_G \max\{c - \omega(\hat{c}), 0\} - \kappa \max\{\hat{c} - \overline{c}, 0\} \tag{5}$$

where $\kappa \max\{\hat{c} - \overline{c}, 0\}$, with $\kappa > 0$, is a penalty suffered for cheating (i.e. presenting a curriculum vitae quality higher than the actual one $\overline{c}$). Note that no penalty is suffered for presenting a curriculum vitae of lower quality than $\overline{c}$.

A worker thus chooses to present a curriculum vitae of quality $\hat{c}$ such that

$$\hat{c} \in \underset{\tilde{c} \in [0,1]}{\operatorname{argmax}}\ u_G(\tilde{c}, c)$$

Definition 3 *Given a wage function $\omega : [0, 1] \to [0, 1]$, a curriculum vitae declaration function $\mu : [0, 1] \to [0, 1]$ assigns a* declared *curriculum vitae quality $\hat{c}$ to an actual curriculum vitae quality $\overline{c}$, that is $\hat{c} = \mu(\overline{c})$.*

Definition 4 *Given a government policy strategy σ^*, a labor market equilibrium (ω^*, μ^*) is a continuous wage function and a curriculum vitae declaration function such that*

$$\omega^*(\hat{c}) = \mathbb{E}[c|\hat{c}, \mu^*, \sigma^*]$$

and

$$\mu^*(\overline{c}) \in \underset{\tilde{c} \in [0,\overline{c}]}{\operatorname{argmax}}\ u_G(\tilde{c}, c).$$

Recall from Eq.(5) that the utility $u_G(\hat{c}, c)$ depends on the wage $\omega^*(\hat{c})$.

If κ is high enough, a continuous wage function $\omega(\hat{c})$ will prevent cheating since the marginal penalty of presenting a curriculum vitae quality greater than $\overline{c}$ will exceed the marginal benefit in terms of increased wage. A sufficient condition for this to hold is that $\kappa > \frac{\omega(\hat{c}) - \omega(\overline{c})}{\hat{c} - \overline{c}}$ for any $\hat{c} > \overline{c}$.

We thus have the following lemma.

Lemma 3 *Suppose κ is high enough. Given a policy strategy σ^*, there exist intervals $\{(c_l^L, c_l^H)\}_{l=1}^{\bar{l}}$ with $\bar{l} \geq 0$, so that the (weakly) increasing wage function*

$$\omega^*(\hat{c}) = \begin{cases} \mathbb{E}[c|\bar{c} = \hat{c}, \sigma^*] & \text{if } \hat{c} \notin \bigcup_l (c_l^L, c_l^H) \\ \mathbb{E}[c|\bar{c} \in (c_l^L, c_l^H), \sigma^*] & \text{if } \hat{c} \in (c_l^L, c_l^H) \end{cases} \tag{6}$$

and the curriculum vitae declaration strategy

$$\mu^*(\bar{c}) = \begin{cases} \bar{c} & \text{if } \bar{c} \notin \bigcup_l (c_l^L, c_l^H) \\ c_l^L & \text{if } \bar{c} \in (c_l^L, c_l^H) \end{cases} \tag{7}$$

constitute a labor market equilibrium.

In the above,

$$\mathbb{E}[c|\bar{c} = \hat{c}, \sigma^*] = \sum_{t=1}^{\infty} \mathbb{P}(t) \left(\mathbb{P}_t^*(\{aa\}|\bar{c}) \cdot g^{-1}(\bar{c}) + \left(1 - \mathbb{P}_t^*(\{aa\}|\bar{c})\right) \cdot \bar{c} \right),$$

with

$$\mathbb{P}_t^*(\{aa\}|\bar{c}) = \frac{|B|\sigma_t^* f_{B,n_t}(g^{-1}(\bar{c}))/g'^{-1}(\bar{c})}{|A|f_A(\bar{c}) + |B|(1-\sigma_t^*)f_{B,n_t}(\bar{c}) + |B|\sigma_t^* f_{B,n_t}(g^{-1}(\bar{c}))/g'^{-1}(\bar{c})},$$

$\{aa\}$ being the event that a worker benefited from affirmative action and

$$\mathbb{P}(t) = \delta^{t-1}(1-\delta)$$

is the probability of being at time t, while

$$\mathbb{E}[c|\bar{c} \in (c_l^L, c_l^H), \sigma^*] = \sum_{t=1}^{\infty} \mathbb{P}(t) \int_{\bar{c}=c_l^L}^{c_l^H} \mathbb{E}[c|\bar{c} = \hat{c}, \sigma^*] f_t(\bar{c}) d\bar{c}$$

where

$$f_t(\bar{c}) = \frac{1}{|A| + |B|} \Big(|A|f_A(\bar{c}) + |B|\sigma_t^* f_{B,n_t}(g^{-1}(\bar{c}))/g'^{-1}(\bar{c}) + |B|(1-\sigma_t^*)f_{B,n_t}(\bar{c}) \Big)$$

is the overall population density for the curriculum vitae quality at time t.

The equilibrium wage function stated in Lemma 3 has the form described in Figure 2(a). We see that it is weakly increasing, but strictly increasing in certain sections. In the particular case when $\bar{l} = 0$, then it can be strictly increasing over the whole domain, as in the case presented earlier in the main part of the paper. The equilibrium curriculum vitae declaration function has the form described in Figure 2(b). It is such that a worker truthfully declares his curriculum vitae quality, i.e. $\hat{c} = \bar{c}$, when $\bar{c}$ is in an interval where the wage function is strictly increasing, since declaring anything lower would yield a lower salary. On the other hand, when $\bar{c}$ is in an interval where the wage function is flat, the worker declares the lowest curriculum vitae quality $\hat{c}$ on that flat interval, i.e. $\hat{c} = c_l^L$. Indeed, declaring such a curriculum vitae quality $\hat{c} \leq \bar{c}$ provides the worker with the same salary as he would get when declaring the actual one: $\omega(\hat{c}) = \omega(\bar{c})$. In the particular case where $\bar{l} = 0$ and the wage function is strictly increasing, then all workers would always declare their true curriculum vitae quality ($\mu^*(\bar{c}) = \bar{c}$, as in the case presented earlier in the main part of the

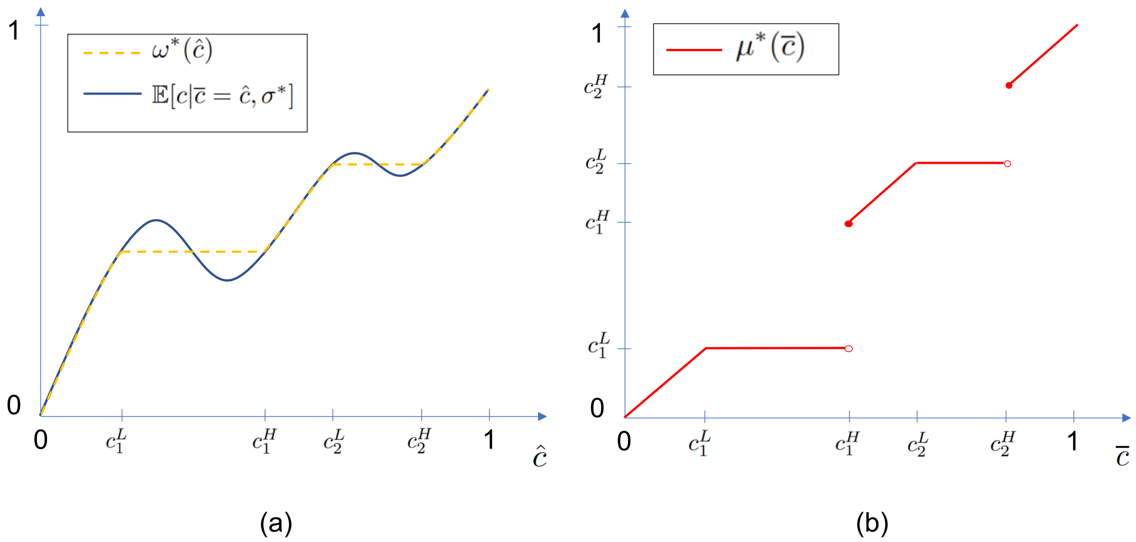

Figure 2: Equilibrium wage function ω^* (panel (a)) and curriculum vitae declaration function μ^* (panel (b)).

paper). Note that, as required by the equilibrium definition, an employer correctly sets the wage equal to the conditional expectation of a worker's performance (i.e. $\omega^*(\hat{c}) = \mathbb{E}[c|\hat{c}, \mu^*, \sigma^*]$).

Lemma 4 *The equilibrium wage function $\omega^*(\hat{c})$ is weakly increasing, but strictly increasing at least on some regions of the support*[21] *$[0, 1]$.*

The next lemma is simply a more general version of Lemma 2, adapted to the labor market equilibrium concept defined in Definition 4.

Lemma 5 *If a worker benefits from affirmative action (i.e. $c = g^{-1}(\overline{c})$), then he gets a wage higher than his performance level (i.e. $c < \omega^*(\mu^*(\overline{c}))$). If a worker does not benefit from affirmative action (i.e. $c = \overline{c}$), then he gets a wage lower than his performance level (i.e. $\omega^*(\mu^*(\overline{c})) < c$).*

Using Lemma 5, we can make the same observations as in Observations 1, 2 and 3, namely that non-beneficiaries of affirmative action (of either group A or B) suffer a feeling of injustice, while beneficiaries do not.

6.3 Proofs

Proof of Lemma 1.

Lemma 1 is a corollary of Lemma 3 when $\mathbb{E}[c|\hat{c} = \overline{c}, \sigma^*]$ is increasing (meaning $\overline{l} = 0$) so that a worker truthfully declares his curriculum vitae quality $\overline{c}$. In such a case, $\omega^*(\hat{c}) = \mathbb{E}[c|\hat{c}, \mu^*, \sigma^*]$ and $\mu^*(\overline{c}) = \overline{c}$ constitute a labor market equilibrium. ■

Proof of Lemma 2.

Lemma 2 is a corollary of Lemma 5 when $\mathbb{E}[c|\hat{c} = \overline{c}, \sigma^*]$ is increasing (meaning $\overline{l} = 0$) so that a worker truthfully declares his curriculum vitae quality $\overline{c}$. In such a case, $\omega^*(\hat{c}) = \mathbb{E}[c|\hat{c}, \mu^*, \sigma^*]$ and $\mu^*(\overline{c}) = \overline{c}$ constitute a labor market equilibrium. ■

[21] This is stronger than needed. $\omega^*(\hat{c})$ only needs to have these properties for the $\hat{c}$'s being played in equilibrium (i.e. $\hat{c} = \mu^*(\overline{c})$).

Proof of Lemma 3.

Throughout this proof, we suppose κ is high enough to prevent cheating. A sufficient condition for this to hold is that $\kappa > \frac{\omega(\hat{c})-\omega(\bar{c})}{\hat{c}-\bar{c}}$ for any $\hat{c} > \bar{c}$. In such a case, the marginal penalty of presenting a curriculum vitae quality greater than $\bar{c}$ will exceed the marginal benefit in terms of increased wage.

Step I: Compute the wage $\tilde{\omega}$ assuming truthful declaration of $\bar{c}$.

Suppose first that workers truthfully declare their curriculum vitae quality, i.e. $\hat{c} = \mu(\bar{c}) = \bar{c}$. Under such a declaration function μ, call $\tilde{\omega}(\hat{c}) = \mathbb{E}[c|\hat{c}, \mu, \sigma^*]$ the conditional expectation of the actual performance level when declaring a curriculum vitae of quality $\hat{c}$. Then,

$$\begin{aligned}
\tilde{\omega}(\hat{c}) &= \mathbb{E}[c|\hat{c}, \mu, \sigma^*] \\
&= \mathbb{E}[c|\hat{c} = \bar{c}, \sigma^*] \\
&= \mathbb{E}\big[\mathbb{E}[c|\hat{c} = \bar{c}, \sigma^*, t]|\hat{c} = \bar{c}, \sigma^*\big] \\
&= \sum_{t=1}^{\infty} \mathbb{P}(t)\mathbb{E}[c|\hat{c} = \bar{c}, \sigma^*, t] \\
&= \sum_{t=1}^{\infty} \mathbb{P}(t)\left(\mathbb{P}_t^*(\{aa\}|\bar{c}) \cdot g^{-1}(\bar{c}) + \big(1 - \mathbb{P}_t^*(\{aa\}|\bar{c})\big) \cdot \bar{c}\right)
\end{aligned}$$

where $\mathbb{P}(t)$ is the probability of being at time t.

Now to express $\mathbb{P}_t^*(\{aa\}|\bar{c})$, we first express $\mathbb{P}_t^*(\{aa\}|\tilde{c} \in N(\bar{c}, \epsilon))$, where $N(\bar{c}, \epsilon)$ is an ϵ-neighborhood of $\bar{c}$:

$$\begin{aligned}
\mathbb{P}_t^*(\{aa\}|\tilde{c} \in N(\bar{c}, \epsilon)) &= \frac{\mathbb{P}_t^*(\{\tilde{c} \in N(\bar{c}, \epsilon)\} \bigcap \{aa\})}{\mathbb{P}_t^*(\tilde{c} \in N(\bar{c}, \epsilon))} \\
&= \frac{\mathbb{P}_t^*(\tilde{c} \in N(\bar{c}, \epsilon) \bigcap B) \cdot \sigma_t^*}{\mathbb{P}_t^*(\tilde{c} \in N(\bar{c}, \epsilon))} \\
&= \frac{\mathbb{P}_t^*(\tilde{c} \in N(\bar{c}, \epsilon)|B) \cdot \mathbb{P}(B) \cdot \sigma_t^*}{\mathbb{P}_t^*(\tilde{c} \in N(\bar{c}, \epsilon))}
\end{aligned}$$

$$\begin{aligned}
&= \frac{\sigma_t^* \int_{\tilde{c} \in N(g^{-1}(\bar{c}), \epsilon/g'^{-1}(\bar{c}))} f_{B,n_t}(g^{-1}(\tilde{c}))d\tilde{c}\frac{|B|}{|A|+|B|}}{\int_{\tilde{c} \in N(\bar{c}, \epsilon)} f_A(\tilde{c})d\tilde{c}\frac{|A|}{|A|+|B|} + (1-\sigma_t^*)\int_{\tilde{c} \in N(\bar{c}, \epsilon)} f_{B,n_t}(\tilde{c})d\tilde{c}\frac{|B|}{|A|+|B|} + \sigma_t^* \int_{\tilde{c} \in N(g^{-1}(\bar{c}), \epsilon/g'^{-1}(\bar{c}))} f_{B,n_t}(g^{-1}(\tilde{c}))d\tilde{c}\frac{|B|}{|A|+|B|}} \\
&= \frac{|B|\sigma_t^* \int_{\tilde{c} \in N(g^{-1}(\bar{c}), \epsilon/g'^{-1}(\bar{c}))} f_{B,n_t}(g^{-1}(\tilde{c}))d\tilde{c}}{|A|\int_{\tilde{c} \in N(\bar{c}, \epsilon)} f_A(\tilde{c})d\tilde{c} + |B|(1-\sigma_t^*)\int_{\tilde{c} \in N(\bar{c}, \epsilon)} f_{B,n_t}(\tilde{c})d\tilde{c} + |B|\sigma_t^* \int_{\tilde{c} \in N(g^{-1}(\bar{c}), \epsilon/g'^{-1}(\bar{c}))} f_{B,n_t}(g^{-1}(\tilde{c}))d\tilde{c}}
\end{aligned}$$

Then, we take the limit as $\epsilon \to 0$:

$$
\begin{aligned}
\mathbb{P}_t^*(\{aa\}|\tilde{c}) &= \lim_{\epsilon\to 0}\mathbb{P}_t^*(\{aa\}|\tilde{c}\in N(\bar{c},\epsilon)) \\
&= \lim_{\epsilon\to 0}\frac{|B|\sigma_t^* f_{B,n_t}(g^{-1}(\bar{c}))2\epsilon/g'^{-1}(\bar{c})}{|A|f_A(\bar{c})2\epsilon+|B|(1-\sigma_t^*)f_{B,n_t}(\bar{c})2\epsilon+|B|\sigma_t^* f_{B,n_t}(g^{-1}(\bar{c}))2\epsilon/g'^{-1}(\bar{c})} \\
&= \frac{|B|\sigma_t^* f_{B,n_t}(g^{-1}(\bar{c}))/g'^{-1}(\bar{c})}{|A|f_A(\bar{c})+|B|(1-\sigma_t^*)f_{B,n_t}(\bar{c})+|B|\sigma_t^* f_{B,n_t}(g^{-1}(\bar{c}))/g'^{-1}(\bar{c})}
\end{aligned}
$$

We finally show that

$$\mathbb{P}(t)=\delta^{t-1}(1-\delta)$$

is the probability of being at time t.

Let $\tau\geq 1$ be the extinction time of the population and let $p(\tau)=\delta^{\tau-1}(1-\delta)$ be the probability that the population goes extinct at time τ. Then, being at time $t=1$ is compatible with extinction times $\tau\geq 2$, while being at time $t=2$ is compatible with extinction times $\tau\geq 3$, etc. In summary, being at time t is compatible with extinction times $\tau\geq t+1$.

Thus, the probability of passing through time t is the probability that $\tau\geq t+1$:

$$
\begin{aligned}
\tilde{\mathbb{P}}(t) &= \sum_{\tau=t+1}^{\infty} p(\tau) \\
&= \sum_{\tau=t+1}^{\infty} \delta^{\tau-1}(1-\delta) \\
&= (1-\delta)\sum_{\tau=t+1}^{\infty} \delta^{\tau-1} \\
&= (1-\delta)\Big(\sum_{\tau=0}^{\infty}\delta^{\tau}-\sum_{\tau=0}^{t-1}\delta^{\tau}\Big) \\
&= (1-\delta)\Big(\frac{1}{1-\delta}-\frac{1-\delta^t}{1-\delta}\Big) \\
&= \delta^t
\end{aligned}
$$

Employers form a belief about being at time t conditional on being at an information set that contains all decision nodes, i.e. $\{t=1,t=2,t=3,...\}$. We use the notion of consistency proposed by Piccione and Rubinstein (1997) that can receive a frequentist interpretation. That is, letting $\tilde{\mathbb{P}}=\sum_{t=1}^{\infty}\tilde{\mathbb{P}}(t)$, the belief attached by an employer being at t is defined by $\mathbb{P}(t)=\tilde{\mathbb{P}}(t)/\tilde{\mathbb{P}}$ where $\mathbb{P}(t)$ represents the frequency with which an employer is at t (in the unique information set $\{t=1,t=2,....\}$). Formally,

$$
\begin{aligned}
\sum_{t=1}^{\infty}\tilde{\mathbb{P}}(t) &= \sum_{t=1}^{\infty}\delta^t \\
&= \frac{1}{1-\delta}-1 \\
&= \frac{\delta}{1-\delta}
\end{aligned}
$$

and we therefore obtain $\mathbb{P}(t)=\tilde{\mathbb{P}}(t)\frac{1-\delta}{\delta}=\delta^{t-1}(1-\delta)$.

Step II: Such a wage function $\tilde{\omega}$ cannot in general be part of an equilibrium.

Suppose that $\tilde{\omega}$ is increasing for $c \in [0, c_1]$ and decreasing over some interval $[c_1, c_1']$. If the wage function is $\tilde{\omega}$, then a worker with an actual curriculum vitae quality $\overline{c} \in (c_1, c_1']$ will choose to declare a curriculum vitae quality $\hat{c} < \overline{c}$ since he can obtain a higher wage $\tilde{\omega}(\hat{c}) > \tilde{\omega}(\overline{c})$ by doing so. It follows that $\mu(\overline{c}) = \overline{c}$ cannot be part of an equilibrium since $\mu(\overline{c}) \notin \underset{\tilde{c}\in[0,\overline{c}]}{\operatorname{argmax}}\, u_G(\tilde{c}, c)$ for such $\overline{c}$.

Since $\mu(\overline{c}) = \overline{c}$ is not part of an equilibrium, it follows that $\tilde{\omega}(\hat{c}) = \mathbb{E}[c|\overline{c} = \hat{c}, \sigma^*]$ is not equal to the correct conditional expectation $\mathbb{E}[c|\hat{c}, \mu^*, \sigma^*]$ where μ^* is an equilibrium declaration function. Thus, $\tilde{\omega}(\hat{c})$ cannot in general be the equilibrium wage function.

Step III: Building a weakly increasing wage function $\omega^*(\hat{c})$ using $\tilde{\omega}(\hat{c})$.

On the other hand, there exist $c_1^L < c_1$ and $c_1^H \geq c_1'$ such that a wage

$$\omega^*(\hat{c}) = \begin{cases} \tilde{\omega}(\hat{c}), & \text{if } \hat{c} \in [0, c_1^L] \\ \tilde{\omega}(c_1^L) & \text{when } \hat{c} \in (c_1^L, c_1^H] \end{cases} \tag{8}$$

corresponds to $\mathbb{E}[c|\hat{c}, \mu^*, \sigma^*]$, where μ^* is as in the statement of the lemma. Such a pair $\{c_1^L, c_1^H\}$ satisfies

$$\tilde{\omega}(c_l^L) = \sum_{t=1}^{\infty} \mathbb{P}(t) \int_{\overline{c}=c_l^L}^{c_l^H} \tilde{\omega}(\overline{c}) f_t(\overline{c}) d\overline{c} \tag{9}$$

$$\tilde{\omega}(c_l^H) = \sum_{t=1}^{\infty} \mathbb{P}(t) \int_{\overline{c}=c_l^L}^{c_l^H} \tilde{\omega}(\overline{c}) f_t(\overline{c}) d\overline{c} \tag{10}$$

and

$$\sum_{t=1}^{\infty} \mathbb{P}(t) \int_{\overline{c}=c_l^H}^{1} \tilde{\omega}(\overline{c}) f_t(\overline{c}) d\overline{c} > \tilde{\omega}(c_l^H). \tag{11}$$

where

$$f_t(\overline{c}) = \frac{1}{|A| + |B|} \Big(|A| f_A(\overline{c}) + |B| \sigma_t^* f_{B,n_t}(g^{-1}(\overline{c}))/g'^{-1}(\overline{c}) + |B|(1 - \sigma_t^*) f_{B,n_t}(\overline{c}) \Big)$$

is simply the overall population density for the curriculum vitae quality $\overline{c}$ at time t.

By construction, $\omega^*(\hat{c})$ is strictly increasing for $\hat{c} \in [0, c_1^L]$ and flat for $\hat{c} \in (c_1^L, c_1^H]$. This is pictured in Figure 2(a). We will generalize this in Step V below.

Step IV: Verifying that (ω^*, μ^*) is a labor market equilibrium for $\overline{c} \in [0, c_1^H]$.

For any worker with an actual curriculum vitae quality $\overline{c} \in [0, c_1^L]$, the best response to such a wage function is $\mu^*(\overline{c}) = \overline{c} = \underset{\tilde{c}\in[0,\overline{c}]}{\operatorname{argmax}}\, u_G(\tilde{c}, c)$ since $\omega^*(\hat{c})$ is strictly increasing over that range and thus the worker chooses to declare $\hat{c} = \overline{c}$ to maximize his wage. Therefore, $\omega^*(\hat{c}) = \mathbb{E}[c|\hat{c}, \mu^*, \sigma^*] = \mathbb{E}[c|\hat{c} = \overline{c}, \sigma^*] = \tilde{\omega}(\hat{c})$ for $\overline{c} \in [0, c_1^L]$. It follows that ω^* and μ^* satisfy the labor market equilibrium condition for $\overline{c} \in [0, c_1^L]$.

Moreover, for any worker with an actual curriculum vitae quality $\overline{c} \in (c_1^L, c_1^H]$, the best response set to a such a wage function is $[c_1^L, \overline{c}] = \underset{\tilde{c}\in[0,\overline{c}]}{\operatorname{argmax}}\, u_G(\tilde{c}, c)$. A worker is indeed indifferent about declaring any $\hat{c} \in [c_1^L, \overline{c}]$, since it yields a salary $\omega^*(\hat{c}) = \tilde{\omega}(c_1^L)$, which is the maximum the worker can obtain. It follows that $\mu^*(\overline{c}) = c_1^L \in \underset{\tilde{c}\in[0,\overline{c}]}{\operatorname{argmax}}\, u_G(\tilde{c}, c)$. Since, $\omega^*(c_1^L) = \mathbb{E}[c|\hat{c}, \mu^*, \sigma^*] = \mathbb{E}[c|\hat{c} = c_1^L, \mu^*, \sigma^*] = \mathbb{E}[c|\overline{c} \in [c_1^L, c_1^H], \sigma^*] = \tilde{\omega}(c_1^L)$, it follows that ω^* and μ^* satisfies the labor market equilibrium condition for $\overline{c} \in (c_1^L, c_1^H]$.

Step V: Generalizing to $\overline{c} \in [0,1]$.

If $c_1^H < 1$ and $\tilde{\omega}(\overline{c})$ is decreasing over some range(s) in $[c_1^H, 1]$, then an iterative application of conditions (9), (10) and (11) allows to find other pairs $\{c_l^L, c_l^H\}$ such that

$$\omega^*(\hat{c}) = \begin{cases} \mathbb{E}[c|\overline{c} = \hat{c}, \sigma^*] & \text{if } \hat{c} \notin \bigcup_l (c_l^L, c_l^H) \\ \mathbb{E}[c|\overline{c} \in (c_l^L, c_l^H), \sigma^*] & \text{if } \hat{c} \in (c_l^L, c_l^H) \end{cases}$$

and the analysis of Steps II, III and IV generalizes to the rest of the support.

■

Proof of Lemma 4.

This is a corollary of Lemma 3.

Lemma 3 states that $\omega^*(\hat{c}) = \mathbb{E}[c|\overline{c} \in (c_l^L, c_l^H), \sigma^*]$ for any $\hat{c} \in (c_l^L, c_l^H)$, implying that $\omega^*(\hat{c})$ is flat for such $\hat{c}$ (since $\mathbb{E}[c|\overline{c} \in (c_l^L, c_l^H), \sigma^*]$ is a constant).

On the other hand, Lemma 3 states that $\omega^*(\hat{c}) = \mathbb{E}[c|\overline{c} = \hat{c}, \sigma^*]$ when $\hat{c} \notin \bigcup_l (c_l^L, c_l^H)$ and Steps III and V of the proof of Lemma 3 show that $\omega^*(\hat{c})$ is constructed so as to be strictly increasing over such intervals.

■

Proof of Lemma 5.

When $\overline{c} \notin \bigcup_l (c_l^L, c_l^H)$, then from Lemma 3 we know that a worker truthfully declares a curriculum vitae quality $\hat{c} = \overline{c}$ and gets a wage

$$\omega^*(\overline{c}) = \sum_{t=1}^{\infty} \mathbb{P}(t) \left(\mathbb{P}_t^*(\{aa\}|\overline{c}) \cdot g^{-1}(\overline{c}) + \left(1 - \mathbb{P}_t^*(\{aa\}|\overline{c})\right) \cdot \overline{c} \right)$$

Since $g^{-1}(\overline{c}) < \overline{c}$, it follows immediately that $g^{-1}(\overline{c}) < \omega^*(\overline{c}) < \overline{c}$.

Thus, if the worker does not benefit from affirmative action (i.e. $c = \overline{c}$), then $\omega^*(\overline{c}) < c$ and he gets a wage lower than his performance level. On the other hand, if the worker benefits from affirmative action (i.e. $c = g^{-1}(\overline{c})$), then $c < \omega^*(\overline{c})$ and he gets a wage higher than his performance level.

We now show that this is also true when $\overline{c} \in \bigcup_l (c_l^L, c_l^H)$.

Recall from Lemma 3 that the wage function is flat over $[c_l^L, c_l^H]$ and equal to $\omega^*(c_l^L)$. Thus, a worker of performance level $c_l'^L$ who does not benefit from affirmative action gets a wage $\omega^*(c_l^L)$ with $\omega^*(c_l^L) < c'$ and a worker of performance level $c_l''^H$ who does not benefit from affirmative action also gets a wage $\omega^*(c_l^L)$ and $\omega^*(c_l^L) < c''$. Consider now a worker who does not benefit from affirmative action and $\overline{c} \in (c_l^L, c_l^H)$. Then, $c = \overline{c}$ with $c' < c < c''$ and the worker gets a wage $\omega^*(c_l^L)$. It follows that $\omega^*(c_l^L) < c$ and he gets a wage lower than his performance level. This applies to any $\overline{c} \in \bigcup_l (c_l^L, c_l^H)$.

Now again, recall from Lemma 3 that the wage function is flat over $[c_l^L, c_l^H]$ and equal to $\omega^*(c_l^L)$. Thus, a worker of performance level $g^{-1}(c_l^L)$ who benefits from affirmative action gets a wage $\omega^*(c_l^L)$ with $g^{-1}(c_l^L) < \omega^*(c_l^L)$ and a worker of performance level $g^{-1}(c_l^H)$ who benefits from affirmative action also gets a wage $\omega^*(c_l^L)$ and $g^{-1}(c_l^H) < \omega^*(c_l^L)$. Consider now a worker who benefits from affirmative action and $\overline{c} \in (c_l^L, c_l^H)$. Then, $c = g^{-1}(\overline{c})$ with $g^{-1}(c_l^L) < g^{-1}(\overline{c}) < g^{-1}(c_l^H)$ and the worker gets a wage $\omega^*(c_l^L)$. It follows that $c = g^{-1}(\overline{c}) < \omega^*(c_l^L)$ and he gets a wage higher than his performance level. This applies to any $\overline{c} \in \bigcup_l (c_l^L, c_l^H)$. ■

We now prove Proposition 1 for a general shape of $\mathbb{E}[c|\overline{c} = \hat{c}, \sigma^*]$ and using the labor market equilibrium concept of Definition 4 and Lemma 3. The particular case presented in the main part of the paper simply corresponds to a truthful curriculum vitae declaration, i.e. $\mu^*(\overline{c}) = \overline{c}$.

Proof of Proposition 1.

We first have the following simple lemma.

Lemma 6 *Let $h(c)$ be any weakly increasing function that is strictly increasing at least on some opened subinterval of its support $[0,1]$ and is differentiable almost everywhere. If $f \succ \tilde{f}$, where f and $\tilde{f}$ are probability density functions on $[0,1]$ and $\succ$ indicates strict first-order stochastic dominance, then $\int_0^1 h(c)f(c)dc > \int_0^1 h(c)\tilde{f}(c)dc$.*

Proof of Lemma 6.

The inequality rewrites

$$\int_0^1 h(c)\left[f(c) - \tilde{f}(c)\right] dc > 0.$$

After integrating by parts, this can be written as

$$\left[h(c)\left[F(c) - \tilde{F}(c)\right]\right]|_0^1 - \int_0^1 h'(c)\left[F(c) - \tilde{F}(c)\right] dc$$

where F and $\tilde{F}$ are the CDFs associated with the PDFs f and $\tilde{f}$. The first term is equal to 0 since $F(0) = \tilde{F}(0) = 0$ and $F(1) = \tilde{F}(1) = 1$. Moreover, since $h'(c) \geq 0$ almost everywhere with $h'(c) > 0$ on non-trivial parts of the support, the last term is strictly greater than 0 if $F(c) < \tilde{F}(c)$ for all $c \in (0,1)$, i.e. if $f \succ \tilde{f}$. ■

Given some equilibrium decision profile $\sigma^* = \{\sigma_s^*\}_{s=1}^\infty$, any deviation σ_t' at some time t has no impact on the wages at any future time since this deviation is unobserved by employers. Indeed, for any $\tau \geq t$, employers form the wage $\omega^*(\hat{c})$ based on the conditional expectation $\mathbb{E}[c|\hat{c}, \mu^*, \sigma^*]$, which depends on the equilibrium decision profile $\sigma^* = \{\sigma_s^*\}_{s=1}^\infty$, none of these decisions being actually observed. Therefore, $\sum_{s=t}^\infty W_{A,s}\delta^{s-t}$, the discounted future welfare of group A, is completely unaffected by an unobserved deviation to σ_t'. Indeed, recall that $c = \overline{c}$ and thus $\mu^*(\overline{c}) = \mu^*(c)$. Therefore, $W_{A,s} = |A|\int_0^1 u_A(\mu^*(c), c)f_A(c)dc$, where the density function $f_A(c)$ is constant through time and thus not impacted by $\{\sigma_s\}_{s=1}^\infty$, while $u_A(\mu^*(c), c) = \omega^*(\mu^*(c)) - \gamma_A(c - \omega^*(\mu^*(c)))$ and the wage $\omega^*(\mu^*(c))$ is unaffected by an unobserved deviation to σ_t'.

Part (i):

We now show that $\{\sigma_s^*\}_{s=1}^\infty = \{1\}_{s=1}^\infty$ is an equilibrium. Given an equilibrium decision profile $\{\sigma_s^*\}_{s=1}^\infty = \{1\}_{s=1}^\infty$, $\sum_{s=t}^\infty \lambda_B W_{B,s}\delta^{s-t}$ is strictly lower following an unobserved deviation from $\sigma_t^* = 1$ to $\sigma_t' = 0$.

Indeed, at time $s = t$,

$$\begin{aligned} W_{B,t|\sigma'_t=0} &= |B| \int_0^1 u_B(\mu^*(c), c|\{aa\}^c) f_{B,n_t}(c) dc \\ &= |B| \int_0^1 \omega^*(\mu^*(c)) - \gamma_B(c - \omega^*(\mu^*(c))) f_{B,n_t}(c) dc \\ &< |B| \int_0^1 \omega^*(\mu^*(g(c))) f_{B,n_t}(c) dc \\ &= |B| \int_0^1 u_B(\mu^*(g(c)), c|\{aa\}) f_{B,n_t}(c) dc \\ &= W_{B,t|\sigma^*_t=1} \end{aligned}$$

where we have used Observations 2 and 3, the fact that the wage is not affected by an unobserved deviation and the facts that $\omega^*(\mu^*(c)) \leq \omega^*(\mu^*(g(c)))$ and that $c > \omega^*(\mu^*(c))$ by Lemma 5.

Moreover, at times $s > t$, $f_{B,n_s|\sigma'_t=0}(c) \prec f_{B,n_s|\sigma^*_t=1}(c)$ since a deviation to $\sigma'_t = 0$ has the effect of not changing the distribution of performance at time $t+1$ compared to the previous period t. Thus, using Observation 3, and the fact that the wage is not affected by an unobserved deviation, then for all $s > t$,

$$\begin{aligned} W_{B,s|\sigma'_t=0} &= |B| \int_0^1 u_B(\mu^*(g(c)), c|\{aa\}) f_{B,n_s|\sigma'_t=0}(c) dc \\ &= |B| \int_0^1 \omega^*(\mu^*(g(c))) f_{B,n_s|\sigma'_t=0}(c) dc \\ &< |B| \int_0^1 \omega^*(\mu^*(g(c))) f_{B,n_s|\sigma^*_t=1}(c) dc \\ &= |B| \int_0^1 u_B(\mu^*(g(c)), c|\{aa\}) f_{B,n_s|\sigma^*_t=1}(c) dc \\ &= W_{B,s|\sigma^*_t=1} \end{aligned}$$

where the inequality follows from Lemma 6. Indeed, from Lemma 4, $\omega^*(\mu^*(g(c)))$ is strictly increasing in c on parts of its support (and weakly increasing overall).

It follows that as long as $\lambda_B > 0$, then $\sigma^*_t = 1$ for all t will be an equilibrium.

Part (ii):

To show that this is the unique equilibrium, we now have to show that a deviation from $\sigma^*_t = 0$ to $\sigma'_t = 1$ is always desirable for a time-t government. For that purpose, suppose that $\sigma^*_t = 0$ for some t. Then, we must show that $\sum_{s=t}^{\infty} \lambda_B W_{B,s} \delta^{s-t}$ is strictly higher following an unobserved deviation from $\sigma^*_t = 0$ to $\sigma'_t = 1$.

Consider first the effect of this deviation on the welfare at time t of members of group B. The same argument as in Part (i) can be used to show that $W_{B,t|\sigma'_t=1} > W_{B,t|\sigma^*_t=0}$.

Consider now the effect of this deviation on the welfare, at any future time $s > t$, of members of group B who benefit from affirmative action. We know that $f_{B,n_s|\sigma^*_t=0}(c) \prec f_{B,n_s|\sigma'_t=1}(c)$ since a deviation to $\sigma'_t = 1$ has the effect of shifting (in a strict first-order stochastic dominance sense) the distribution of performance from time t to $t+1$. Moreover, since from Observation 3 $u_B(\mu^*(g(c)), c|\{aa\}) = \omega^*(\mu^*(g(c)))$, and the wage is not affected by an unobserved deviation, then for all $s > t$,

$$\begin{aligned} W_{B,s|\{aa\},\sigma_t^*=0} &= |B| \int_0^1 \omega^*(\mu^*(g(c))) f_{B,n_s|\sigma_t^*=0}(c)dc \\ &< |B| \int_0^1 \omega^*(\mu^*(g(c))) f_{B,n_s|\sigma_t'=1}(c)dc \\ &= W_{B,n_s|\{aa\},\sigma_t'=1} \end{aligned}$$

where we made use of Lemma 6, since $\omega^*(\mu^*(g(c)))$ is strictly increasing in c on parts of its support (and weakly increasing overall from Lemma 4) and $f_{B,n_s|\sigma_t^*=0}(c) \prec f_{B,n_s|\sigma_t'=1}(c)$ for all $s > t$.

Consider finally the effect of this deviation on the welfare, at any future time $s > t$, of members of group B who do *not* benefit from affirmative action. Using Observation 2,

$$\begin{aligned} W_{B,s|\{aa\}^c} &= |B| \int_0^1 \big(\omega^*(\mu^*(c)) - \gamma_B(c - \omega^*(\mu^*(c)))\big) f_{B,n_s}(c)dc \\ &= |B| \int_0^1 \big((1+\gamma_B)\omega^*(\mu^*(c)) - \gamma_B c\big) f_{B,n_s}(c)dc \end{aligned} \tag{12}$$

Moreover, since $f_{B,n_s|\sigma_t'=1}(c) \succ f_{B,n_s|\sigma_t^*=0}(c)$ for all $s > t$ and making use of Lemma 6, note that

$$\int_0^1 \big((1+\gamma_B)\omega^*(\mu^*(c))\big) f_{B,n_s|\sigma_t'=1}(c)dc > \int_0^1 \big((1+\gamma_B)\omega^*(\mu^*(c))\big) f_{B,n_s|\sigma_t^*=0}(c)dc$$

for all $\gamma_B \geq 0$. It thus follows that there exists $\overline{\gamma}_B$ such that for $\gamma_B < \overline{\gamma}_B$,

$$\begin{aligned} &\int_0^1 \big((1+\gamma_B)\omega^*(\mu^*(c))\big) f_{B,n_s|\sigma_t'=1}(c)dc - \int_0^1 \big((1+\gamma_B)\omega^*(\mu^*(c))\big) f_{B,n_s|\sigma_t^*=0}(c)dc \\ &\qquad > \int_0^1 \gamma_B c f_{B,n_s|\sigma_t'=1}(c)dc - \int_0^1 \gamma_B c f_{B,n_s|\sigma_t^*=0}(c)dc \end{aligned} \tag{13}$$

and thus that $W_{B,n_s|\{aa\}^c,\sigma_t'=1} > W_{B,n_s|\{aa\}^c,\sigma_t^*=0}$. We can then conclude that the future welfare of all members of group B (benefitting from affirmative action or not) increases following this unobserved deviation from $\sigma_t^* = 0$ to $\sigma_t' = 1$ and thus that $\sum_{s=t}^{\infty} \lambda_B W_{B,s} \delta^{s-t}$ is indeed strictly higher. Therefore, when the feeling of injustice parameter is small enough, $\gamma_B < \overline{\gamma}_B$, an equilibrium decision profile $\{\sigma_s^*\}_{s=1}^{\infty}$ in which $\sigma_t^* = 0$ for some t cannot exist.

Part (iii):

We will verify that the proof of Part (ii) also holds for any $\gamma_B \geq 0$ as long as $\beta = \frac{|B|}{|A|+|B|}$ is small enough.

First note that $\omega^*(\mu^*(\overline{c}))$ converges to c as $\beta \to 0$. Indeed, as the relative size of the targeted group decreases, the probability that a worker benefited from affirmative action (and thus an inflated curriculum vitae quality) decreases as well and thus the wage becomes equal to the actual performance level c in the limit as β goes to 0.

Thus, for any $\epsilon > 0$, there exists $\overline{\beta}$ such that for all $\beta < \overline{\beta}$, we have $|\int_0^1 \omega^*(\mu^*(\overline{c})) f_{B,n_s}(c)dc - \int_0^1 c f_{B,n_s}(c)dc| < \epsilon$.

Since for any $\gamma_B \geq 0$,

$$(1+\gamma_B)\Big(\int_0^1 cf_{B,n_s|\sigma'_t=1}(c)dc - \int_0^1 cf_{B,n_s|\sigma^*_t=0}(c)dc\Big) > \gamma_B\Big(\int_0^1 cf_{B,n_s|\sigma'_t=1}(c)dc - \int_0^1 cf_{B,n_s|\sigma^*_t=0}(c)dc\Big)$$

then we have that

$$\begin{aligned}(1+\gamma_B)\Big(\int_0^1 \omega^*(\mu^*(\overline{c}))f_{B,n_s|\sigma'_t=1}(c)dc - \int_0^1 \omega^*(\mu^*(\overline{c}))f_{B,n_s|\sigma^*_t=0}(c)dc\Big) &> \\ (1+\gamma_B)\Big(\int_0^1 cf_{B,n_s|\sigma'_t=1}(c)dc - \int_0^1 cf_{B,n_s|\sigma^*_t=0}(c)dc\Big) - 2\epsilon &> \\ \gamma_B\Big(\int_0^1 cf_{B,n_s|\sigma'_t=1}(c)dc - \int_0^1 cf_{B,n_s|\sigma^*_t=0}(c)dc\Big)&\end{aligned}$$

where the last inequality holds for small enough $\epsilon > 0$ and thus when $\beta < \overline{\beta}$.

It then follows that Eq. (13) in Part (ii) also holds for any $\gamma_B \geq 0$ as long as $\beta = \frac{|B|}{|A|+|B|} < \overline{\beta}$. ■

We now prove Proposition 2.

Proof of Proposition 2.

We start with the following lemma.

Lemma 7 *Let $\{\sigma_t\}_{t=1}^{\infty}$ be a policy plan with $\sigma_\tau = 0$ and $\sigma_{\tau+1} = 1$ for some τ. Let $\{\sigma'_t\}_{t=1}^{\infty}$ be another policy plan with $\sigma'_\tau = 1$, $\sigma'_{\tau+1} = 0$ and $\sigma'_t = \sigma_t$ for all other t. Then there exists $\overline{\delta} \geq 0$ such that for all $\delta \in (\overline{\delta}, 1)$, $\{\sigma'_t\}_{t=1}^{\infty}$ yields a strictly higher welfare than $\{\sigma_t\}_{t=1}^{\infty}$.*

Proof of Lemma 7. We will prove this lemma for a general shape of $\mathbb{E}[c|\overline{c} = \hat{c}, \sigma^*]$ and using the labor market equilibrium concept of Definition 4 and Lemma 3.

Suppose first that $\delta = 1$.

Then $\{\sigma'_t\}_{t=1}^{\infty}$ and $\{\sigma_t\}_{t=1}^{\infty}$ yield the same wage function[22] $\omega^*(\hat{c})$ since $\mathbb{P}\{\tau\} = \mathbb{P}\{\tau+1\}$.

Therefore $\sum_{t=1}^{\infty}\delta^t W_{A,t} = \sum_{t=1}^{\infty}\delta^t|A|\int_0^1\big((1+\gamma_A)\omega^*(\mu^*(c)) - \gamma_A c\big)f_A(c)dc$ is the same under plans $\{\sigma'_t\}_{t=1}^{\infty}$ and $\{\sigma_t\}_{t=1}^{\infty}$.

On the other hand, $\lambda_B\sum_{t=1}^{\infty}\delta^t W_{B,t}$ is strictly greater under plan $\{\sigma'_t\}_{t=1}^{\infty}$ than under $\{\sigma_t\}_{t=1}^{\infty}$. To see this, note that $\lambda_B\sum_{t=1}^{\tau-1}\delta^t W_{B,t}$ is the same under both policy plans, while $\lambda_B\sum_{t=\tau}^{\infty}\delta^t W_{B,t} = \lambda_B\sum_{t=\tau}^{\infty}\delta^t|B|\int_0^1 u_B(\mu^*(\overline{c}), c)f_{B,n_t}(c)dc$ is strictly greater under plan $\{\sigma'_t\}_{t=1}^{\infty}$ than under $\{\sigma_t\}_{t=1}^{\infty}$. Indeed, under $\{\sigma'_t\}_{t=1}^{\infty}$,

$$\sum_{t=\tau}^{\infty}\delta^t W'_{B,t} = \delta^\tau|B|\int_0^1 \omega^*(\mu^*(g(c)))f_{B,n'\tau}(c)dc + \delta^{\tau+1}|B|\int_0^1\big((1+\gamma_B)\omega^*(\mu^*(c)) - \gamma_B c\big)f_{B,n'_{\tau+1}}(c)dc + \sum_{t=\tau+2}^{\infty}\delta^t W'_{B,t}$$

while under $\{\sigma_t\}_{t=1}^{\infty}$

$$\sum_{t=\tau}^{\infty}\delta^t W_{B,t} = \delta^\tau|B|\int_0^1\big((1+\gamma_B)\omega^*(\mu^*(c)) - \gamma_B c\big)f_{B,n\tau}(c)dc + \delta^{\tau+1}|B|\int_0^1 \omega^*(\mu^*(g(c)))f_{B,n_{\tau+1}}(c)dc + \sum_{t=\tau+2}^{\infty}\delta^t W_{B,t}$$

[22] Formally, we would take a limit as $\delta \to 1$ and note that $\omega'^*(\hat{c}) \xrightarrow[\delta\to 1]{} \omega^*(\hat{c})$.

Only the first two terms differ and their sum is greater under $\{\sigma'_t\}_{t=1}^{\infty}$ $\underline{\text{when } \delta = 1}$. Indeed, this follows from the facts that $\omega^*(\mu^*(g(c))) > \big((1+\gamma_B)\omega^*(\mu^*(c)) - \gamma_B c\big)$, that $f_{B,n'_{\tau+1}}(c) \succ f_{B,n\tau}(c)$ and that $f_{B,n'\tau}(c) = f_{B,n_{\tau+1}}(c)$.

By continuity, it follows that there exists[23] $\overline{\delta} \in (0,1)$ such that for all $\delta \in (\overline{\delta}, 1)$, the total welfare is also higher under plan $\{\sigma'_t\}_{t=1}^{\infty}$ than under $\{\sigma_t\}_{t=1}^{\infty}$.

■

Therefore, when δ is high enough, it follows by iterative application of Lemma 7 that the optimal policy has a threshold form $\hat{\sigma}_t = 1$ for $t < \overline{T}$ and $\hat{\sigma}_t = 0$ for $t \geq \overline{T}$ for some $\overline{T} \in \mathbb{N} \bigcup \infty$.

We will now rule out the case where $\overline{T}$ could be infinite and thus show that $\overline{T} \in \mathbb{N}$.

To reduce the notational burden, we will prove this in the setting where $\mathbb{E}[c|\overline{c} = \hat{c}, \sigma^*]$ in non-decreasing (and thus $\mu^*(\overline{c}) = \overline{c}$), as in the main part of the paper. Naturally, the result still holds using the labor market equilibrium concept of Definition 4 and Lemma 3.

Let us thus compare the welfare of some (large) $\overline{T} < \infty$ to that of the case $\overline{T}' = \infty$. In what follows, the quantities with a prime ($'$) will be the ones associated to $\overline{T}' = \infty$.

We need to show that

$$\sum_{t=1}^{\infty} \delta^t (W_{A,t} + \lambda_B W_{B,t}) > \sum_{t=1}^{\infty} \delta^t (W'_{A,t} + \lambda_B W'_{B,t}). \tag{14}$$

Equivalently, it will be convenient to multiply the welfare by the constant $\frac{(1-\delta)}{\delta(|A|+|B|)}$ and verify that

$$\frac{(1-\delta)}{\delta(|A|+|B|)} \Big(\sum_{t=1}^{\infty} \delta^t (W_{A,t} + \lambda_B W_{Bt}) - \sum_{t=1}^{\infty} \delta^t (W'_{A,t} + \lambda_B W'_{B,t}) \Big) > 0$$

First recall that $\mathbb{P}(t) = \delta^{t-1}(1-\delta)$ and we can rewrite the left-hand side of this expression as

$$\begin{aligned}
\sum_{t=1}^{\infty} \frac{\delta^{t-1}(1-\delta)}{(|A|+|B|)} \big((W_{A,t} + \lambda_B W_{B,t}) - (W'_{A,t} + \lambda_B W'_{B,t})\big) &= \sum_{t=1}^{\infty} \frac{\mathbb{P}(t)}{(|A|+|B|)} \big((W_{A,t} + \lambda_B W_{B,t}) - (W'_{A,t} + \lambda_B W'_{B,t})\big) \\
&= \sum_{t=1}^{\infty} \mathbb{P}(t) \Big(\frac{|A|}{|A|+|B|} \int \omega^*(c) f_A(c) dc \\
&\quad + \frac{\lambda_B |B|}{|A|+|B|} \int [\sigma_t \omega^*(g(c)) + (1-\sigma_t)\omega^*(c)] f_{B,n_t}(c) dc \\
&\quad - \frac{|A|}{|A|+|B|} \int \omega'^*(c) f_A(c) dc \\
&\quad - \frac{\lambda_B |B|}{|A|+|B|} \int \omega'^*(g(c)) f_{B,n_t}(c) dc \Big) \\
&\quad + \sum_{t=1}^{\infty} \mathbb{P}(t) \frac{|A|}{|A|+|B|} \gamma_A \int (\omega^*(c) - \omega'^*(c)) f_A(c) dc \\
&\quad - \sum_{t=\overline{T}}^{\infty} \mathbb{P}(t) \frac{\lambda_B |B|}{|A|+|B|} \gamma_B \int (c - \omega^*(c)) f_{B,n_{\overline{T}}}(c) dc
\end{aligned} \tag{15}$$

The case $\lambda_B = 1$ is interesting and worth examining first. In that case, note that the first two

[23] Note that even if $f_{B,n'_{\tau+1}}(c)$ is arbitrarily close to $f_{B,n_\tau}(c)$ (which will happen when n_τ is large), the argument goes through as long as $\omega^*(\mu^*(g(c))) > (1+\gamma_B)\omega^*(\mu^*(c)) - \gamma_B c$ (which follows from Lemmas 4 and 5). This means δ is not required to get closer and closer to 1 as τ increases to infinity.

terms of the right-hand side of Eq. (15) rewrite as

$$\sum_{t=1}^{\infty}\mathbb{P}(t)\Big(\frac{|A|}{|A|+|B|}\int \omega^*(c)f_A(c)dc+\frac{|B|}{|A|+|B|}\int[\sigma_t\omega^*(g(c))+(1-\sigma_t)\omega^*(c)]f_{B,n_t}(c)dc\Big)=\mathbb{E}[c], \tag{16}$$

since the time-wise average wage under policy $\overline{T}<\infty$ is equal to the time-wise average performance level under policy $\overline{T}<\infty$ (here denoted by $\mathbb{E}[c]$).

Likewise, the third and fourth terms rewrite as

$$-\sum_{t=1}^{\infty}\mathbb{P}(t)\Big(\frac{|A|}{|A|+|B|}\int \omega'^*(c)f_A(c)dc+\frac{|B|}{|A|+|B|}\int \omega'^*(g(c))f_{B,n_t}(c)dc\Big)=-\mathbb{E}'[c], \tag{17}$$

since the time-wise average wage under policy $\overline{T}'=\infty$ is equal to the time-wise average performance level under policy $\overline{T}'=\infty$ (here denoted by $\mathbb{E}'[c]$).

We then have that the right-hand side of Eq. (15) can be written as

$$\big(\mathbb{E}[c]-\mathbb{E}'[c]\big)+\frac{|A|}{|A|+|B|}\gamma_A\int(\omega^*(c)-\omega'^*(c))f_A(c)dc-\frac{|B|}{|A|+|B|}\gamma_B\int(c-\omega^*(c))f_{B,n_{\overline{T}}}(c)dc\sum_{t=\overline{T}}^{\infty}\mathbb{P}(t)$$

where we have used the fact that $\sum_{t=1}^{\infty}\mathbb{P}(t)=1$ in the fifth term of Eq. (15).

We must now verify if this is greater than 0. We first make the following observations:

- The first term, $\mathbb{E}[c]-\mathbb{E}'[c]<0$. Indeed, the time-wise average performance level is higher when affirmative action is implemented for more periods. Moreover, this term converges to 0 as $\overline{T}\to\infty$ since $\mathbb{E}[c]\to\mathbb{E}'[c]$, capturing the fact that the improvements in the performance distribution of group B become marginal after a while.

- The third term is smaller than 0 since $c-\omega^*(c)>0$ as the members of group B suffer a feeling of injustice at times greater or equal to $\overline{T}$, when affirmative action has been stopped. On the other hand, for any policy $\overline{T}<\infty$, this term also converges to 0 as $\delta\to 1$. Indeed when affirmative action is stopped after a finite number of periods, the probability that a worker presenting a curriculum vitae c benefited from affirmative action goes to 0 when the survival probability becomes large enough, i.e. $\mathbb{P}^*(\{aa\}|c)=\sum_{t=1}^{\infty}\mathbb{P}(t)\mathbb{P}_t^*(\{aa\}|c)\xrightarrow[\delta\to 1]{}0$. It then follows that in such a case, a worker is paid his actual performance level, i.e. $\omega^*(c)\xrightarrow[\delta\to 1]{}c$. Moreover, since $\sum_{s=\overline{T}}^{\infty}\mathbb{P}(t)\leq 1$, it follows that the third term converges to 0 as $\delta\to 1$. The interpretation is that the feeling of injustice becomes negligible in such a case under policy $\overline{T}<\infty$.

- The second term is positive since $\omega^*(c)-\omega'^*(c)>0$. Moreover, under a policy of permanent affirmative action $\overline{T}'=\infty$,

$$\begin{aligned}\omega'^*(c) &= \sum_{t=1}^{\infty}\mathbb{P}(t)\big(\mathbb{P}_t'^*(\{aa\}|c)g^{-1}(c)+(1-\mathbb{P}_t'^*(\{aa\}|c))c\big)\\ &< c\end{aligned}$$

  where the second equality comes from the fact that $\mathbb{P}_t'^*(\{aa\}|c)>0$ for all t. Since as seen previously, $\omega^*(c)\xrightarrow[\delta\to 1]{}c$ for any $\overline{T}<\infty$, then we conclude that $\omega^*(c)-\omega'^*(c)\xrightarrow[\delta\to 1]{}\Delta$ for some $\Delta > 0$ and thus the second term is bounded away from 0 in the limit. This captures the gain to group A of stopping affirmative action after a finite number of periods.

From the above three observations, we can formally state that $\forall \epsilon > 0$, there exists $\overline{T} < \infty$ large enough and $\overline{\delta}(\overline{T}) \in (0,1)$ such that $\forall \delta \in (\overline{\delta}(\overline{T}), 1)$

$$|\mathbb{E}[c] - \mathbb{E}'[c]| < \epsilon,$$

$$\frac{|B|}{|A|+|B|}\gamma_B \int (c - \omega^*(c)) f_{B,n_{\overline{T}}}(c)dc \sum_{s=\overline{T}}^{\infty} \mathbb{P}(t) < \epsilon$$

and

$$\frac{|A|}{|A|+|B|}\gamma_A \int (\omega^*(c) - \omega'^*(c)) f_A(c)dc > 2\epsilon$$

from which it follows that the right-hand side of Eq. (15) is positive and thus that Eq. (14) is verified.

To complete the proof, we now turn to the case when $\lambda_B < 1$.

First note that, unsurprisingly, group A gains from stopping affirmative action, whereas group B loses (when δ is high enough). Thus, rearranging the left-hand side of Eq. (15) as follows

$$\sum_{t=1}^{\infty} \frac{\delta^{t-1}(1-\delta)}{(|A|+|B|)} \big((W_{A,t} - W'_{A,t}) + \lambda_B(W_{B,t}) - W'_{B,t})\big),$$

we notice that decreasing the weight λ_B placed on the welfare of group B to values strictly smaller than 1 keeps this quantity positive. We can thus conclude that it will still be worth stopping affirmative action after $\overline{T} < \infty$ periods as opposed to continuing it forever. The first-best optimal policy $\overline{T}_{\lambda_B}$ for some $\lambda_B < 1$ will thus be such that $\overline{T}_{\lambda_B} \leq \overline{T}_{\lambda_B=1} < \infty$.

Let us now verify that, group A gains from stopping affirmative action, whereas group B loses. From the decomposition in Eq. (15), we can write the welfare gain of group A from stopping affirmative action as

$$\sum_{t=1}^{\infty} \frac{\delta^{t-1}(1-\delta)}{(|A|+|B|)} \big(W_{A,t} - W'_{A,t}\big) = \sum_{t=1}^{\infty} \mathbb{P}(t) \frac{|A|}{|A|+|B|}(1+\gamma_A) \int (\omega^*(c) - \omega'^*(c)) f_A(c)dc > 0.$$

This is positive, even when $\gamma_A = 0$, since $\omega^*(c) > \omega'^*(c)$. We can also write the welfare gain associated with group B from stopping affirmative action as

$$\begin{aligned}\sum_{t=1}^{\infty} \frac{\delta^{t-1}(1-\delta)\lambda_B}{(|A|+|B|)} \big(W_{B,t} - W'_{B,t}\big) \quad &= \quad \sum_{t=1}^{\infty} \mathbb{P}(t) \frac{\lambda_B|B|}{|A|+|B|} \int [\sigma_t \omega^*(g(c)) + (1-\sigma_t)\omega^*(c)] f_{B,n_t}(c)dc \\ &\quad - \sum_{t=\overline{T}}^{\infty} \mathbb{P}(t) \frac{\lambda_B|B|}{|A|+|B|}\gamma_B \int (c - \omega^*(c)) f_{B,n_{\overline{T}}}(c)dc \\ &\quad - \sum_{t=1}^{\infty} \mathbb{P}(t) \frac{\lambda_B|B|}{|A|+|B|} \int \omega'^*(g(c)) f_{B,n_t}(c)dc. \end{aligned} \tag{18}$$

The second term

$$-\sum_{t=\overline{T}}^{\infty} \mathbb{P}(t) \frac{\lambda_B|B|}{|A|+|B|}\gamma_B \int (c - \omega^*(c)) f_{B,n_{\overline{T}}}(c)dc \xrightarrow[\delta \to 1]{} 0$$

since $\omega^*(c) \xrightarrow[\delta \to 1]{} c$ and is negative for any $\delta < 1$ since $c > \omega^*(c)$, i.e. members of group B suffer a feeling of injustice after affirmative action has been stopped.

The first term can be split into

$$\sum_{t=1}^{\infty} \mathbb{P}(t) \frac{\lambda_B |B|}{|A| + |B|} \int \sigma_t \omega^*(g(c)) f_{B,n_t}(c) dc = \sum_{t=1}^{\overline{T}-1} \mathbb{P}(t) \frac{\lambda_B |B|}{|A| + |B|} \int \omega^*(g(c)) f_{B,n_t}(c) dc \xrightarrow[\delta \to 1]{} 0$$

and

$$\begin{aligned} \sum_{t=1}^{\infty} \mathbb{P}(t) \frac{\lambda_B |B|}{|A| + |B|} \int (1 - \sigma_t) \omega^*(c) f_{B,n_t}(c) dc &= \frac{\lambda_B |B|}{|A| + |B|} \int \omega^*(c) f_{B,n_{\overline{T}}}(c) dc \sum_{t=\overline{T}}^{\infty} \mathbb{P}(t) \\ &\xrightarrow[\delta \to 1]{} \frac{\lambda_B |B|}{|A| + |B|} \int c f_{B,n_{\overline{T}}}(c) dc \sum_{t=\overline{T}}^{\infty} \mathbb{P}(t) \\ &> 0 \end{aligned}$$

since $\omega^*(c) \xrightarrow[\delta \to 1]{} c$.

Finally, the third term

$$-\sum_{t=1}^{\infty} \mathbb{P}(t) \frac{\lambda_B |B|}{|A| + |B|} \int \omega'^*(g(c)) f_{B,n_t}(c) dc < 0$$

is bounded below 0.

Thus as δ approaches 1, the right hand side of Eq. (18) approaches

$$\frac{\lambda_B |B|}{|A| + |B|} \int c f_{B,n_{\overline{T}}}(c) dc \sum_{t=\overline{T}}^{\infty} \mathbb{P}(t) - \sum_{t=1}^{\infty} \mathbb{P}(t) \frac{\lambda_B |B|}{|A| + |B|} \int \omega'^*(g(c)) f_{B,n_t}(c) dc < 0$$

which is indeed negative since $\omega'^*(g(c)) > c$ (i.e. beneficiaries of affirmative action do not suffer a feeling of injustice) and since f_{B,n_t} first-order stochastically dominates $f_{B,n_{\overline{T}}}$ for $t > \overline{T}$. Group B would therefore benefit from affirmative action continuing permanently.

Note that when $\lambda_B < 1$, the parameter γ_A governing the feeling of injustice of group A can be 0 and the first-best policy would still prescribe stopping affirmative action after a finite number of periods. On the other hand, if λ_B were to be strictly greater than 1 and thus if the government placed more weight on the welfare of group B than that of group A, then it might become necessary at some point for γ_A to be sufficiently positive in order to justify stopping affirmative action. ■

References

Arrow, K. (1973): "The theory of discrimination. In: Discrimination in labor markets," *Orley Achenfelter and Albert Rees (eds.), Princeton–New Jersey.*

Bagwell, K. (1995): "Commitment and observability in games," *Games and Economic Behavior*, 8, 271–280.

Becker, G. S. (1957): *The economics of discrimination*, University of Chicago press.

Chung, K.-S. (2000): "Role models and arguments for affirmative action," *American Economic Review*, 90, 640–648.

Coate, S. and G. Loury (1993a): "Antidiscrimination enforcement and the problem of patronization," *American Economic Review*, 83, 92–98.

Coate, S. and G. C. Loury (1993b): "Will affirmative-action policies eliminate negative stereotypes?" *American Economic Review*, 1220–1240.

Fang, H. and A. Moro (2011): "Theories of statistical discrimination and affirmative action: A survey," *Handbook of social economics*, 1, 133–200.

Kahneman, D. and A. Tversky (1979): "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, 47, 263–292.

Lundberg, S. J. and R. Startz (1983): "Private discrimination and social intervention in competitive labor market," *American Economic Review*, 73, 340–347.

Moro, A. and P. Norman (2003): "Affirmative action in a competitive economy," *Journal of Public Economics*, 87, 567–594.

Phelps, E. S. (1972): "The statistical theory of racism and sexism," *American Economic Review*, 62, 659–661.

Piccione, M. and A. Rubinstein (1997): "On the interpretation of decision problems with imperfect recall," *Games and Economic Behavior*, 20, 3–24.

Van Damme, E. and S. Hurkens (1997): "Games with imperfectly observable commitment," *Games and Economic Behavior*, 21, 282–308.