Journal of Economic Behavior & Organization 129 (2016) 116–128

Contents lists available at ScienceDirect

Journal of Economic Behavior & Organization

journal homepage: www.elsevier.com/locate/jebo

From taste-based to statistical discrimination☆

William Neilson[a,*], Shanshan Ying[b]

[a] *Department of Economics, University of Tennessee—Knoxville, 508 Stokely Management Center, Knoxville, TN 37919, USA*
[b] *School of International Business and Administration, Shanghai University of Finance and Economics—Shanghai, 777 Guoding Road, Shanghai 200433, China*

ARTICLE INFO

Article history:
Received 5 January 2016
Received in revised form 10 May 2016
Accepted 4 June 2016
Available online 21 June 2016

JEL classifications:
J71
D21
M14

Keywords:
Discrimination
Gender
Beauty
Ascriptive characteristics
Hiring

ABSTRACT

Consider hiring managers who care not just about productivity but also some other, unrelated characteristic. If they treat that ascriptive characteristic differently across groups by, for example, valuing beauty more for women than men, then the hired women will be better looking but less productive, on average. This taste-based discrimination, focused entirely on an ascriptive characteristic, can lead to productivity-based statistical discrimination by the firm's subsequent hiring managers who observe from their workforce that women tend to produce less. This identifies a new channel behind statistical discrimination that arises from the behavior of prior hiring managers.

1. Introduction

The literature on discrimination has led to two different rationales for an employer to favor one group over another. Becker (1957) introduced taste-based discrimination in which a principal simply has a preference for working with one type over the other. Phelps (1972) offered an alternative model in which productivity cannot be observed perfectly and group identity might contain information about productivity, in which case profit maximization would lead to favoring one group over the other. Both models have received empirical support (see, for example, the book by Hamermesh, 2011; and the surveys by Lang and Lehmann, 2012, and Liu and Sierminska, 2014). In this paper we address a different question: can current taste-based discrimination lead to future statistical discrimination?

We answer this question using a theory model, and to see how it works consider the following example. The first hiring manager observes the applicant's gender, ability, and beauty. He does not know whether gender matters for productivity, but he knows that ability does. He also has a taste for beauty, which is unrelated to productivity. Moreover, he cares more about looks for women than he does for men, and trades off looks and ability in hiring decisions. This is taste-based discrimination. We show that this taste leads to a workforce in which women have better looks, on average, than men do, consistent with

☆ This paper was written while Ying visited the University of Tennessee, Knoxville. We are grateful to Yongmin Chen, Dan Hamermesh, Liqun Liu, Marie Claire Villeval, and two reviewers for helpful comments. Errors remain our responsibility.

* Corresponding author.
E-mail addresses: wneilson@utk.edu (W. Neilson), yss37@126.com (S. Ying).

http://dx.doi.org/10.1016/j.jebo.2016.06.001
0167-2681/

the evidence from Hamermesh and Biddle (1994), but men have higher ability, on average, than women. After this first manager leaves the company, his successor could look at the existing workforce to determine whether gender has any relationship to performance. If she did not dig too deeply by controlling for ability, she would infer that men are more productive than women at this firm. If gender can be observed costlessly in the employee search process but ability cannot, profit maximization would then lead to oversampling males when recruiting, which would in turn lead to hiring more males. This is statistical discrimination.

Gender is only one source of group identity of course, and others include race or nationality. Beauty as a basis for discrimination was first investigated by Hamermesh and Biddle (1994), but is not the only ascriptive characteristic that could lead to taste-based discrimination by the first manager. Other researchers have documented discrimination based on such things as skin tone (Hersch, 2008), hair color (Johnston, 2010), height (Hersch, 2008; Persico et al., 2004), weight (Hersch, 2008), or a Southern accent (Kinzler and DeJesus, 2013). The same story works for characteristics as seemingly benign as being a sports fan. If the first hiring manager has a preference for hiring men who can talk with him about sports, but does not care about this attribute for women, his hiring practices would lead to a workforce with men who are relatively more knowledgeable about sports and women who are relatively more productive. The second manager would then find a basis for statistical discrimination against men.

Our main finding is that if the first manager treats any ascriptive characteristic differently across groups when making hiring decisions, the generated workforce will provide a basis for statistical discrimination from subsequent hiring managers. The model is general, with no restrictions on the joint distribution of ability and the ascriptive characteristic other than that they are statistically independent, which is what makes the characteristic ascriptive in the first place. We establish our result using two different models of taste-based discrimination for the first manager. In one model the manager uses the same ability-beauty tradeoff for both genders but has different minimum beauty thresholds for men and women. In the second his marginal rate of substitution between ability and beauty differs for the two groups. In both cases we find that if looks matter more to the manager for female applicants than for men, the hired workforce establishes a basis for statistical discrimination against women.

Our paper adds to the literature on the differences between taste-based and statistical discrimination and the literature on beauty premia. The literature on the sources of discrimination is largely empirical, and those papers have sought to uncover one source or the other (see Guryan and Charles, 2013; for a discussion). Our paper shows that the two sources might be related, with one leading to the other. Work on the beauty premium has also been largely empirical, and our paper provides a theoretical foundation for that literature.

There are several theory papers that demonstrate how the search process can lead to statistical discrimination. In Coate and Loury (1993), employers' negative stereotypes affect the human capital investment decisions of workers, which in turn confirm employers' negative beliefs in equilibrium. Morgan and Vardy (2009) show that when signals of minority candidate ability are noisier than those of majority candidates, minority candidates are less likely to be hired and are therefore underrepresented in the workplace. The paper most like ours is Bagues and Perez-Villadoniga (2013), which also provides a theoretical treatment in which multi-dimensional attributes can lead to statistical discrimination. In our paper one dimension is productive and the other ascriptive, completely orthogonal to productivity. In their paper, in contrast, all dimensions are productive, and statistical discrimination arises when the manager receives more precise signals about one dimension than the others, as in the single-dimensional framework of Morgan and Vardy (2009). The manager then places more weight on that dimension, leading to discrimination in favor of candidates stronger in that dimension and against candidates stronger in the other dimensions.[1] The major difference between our paper and those of Morgan and Vardy (2009) and Bagues and Perez-Villadoniga (2013) is that in their models statistical discrimination arises out of the data-generating process underlying employer search, while in our model statistical discrimination arises out of taste-based discrimination governed by an ascriptive characteristic. We show that the taste-based discrimination can have long-term effects even when the underlying preference bias is removed.

We go on to explore how this channel changes when the underlying wage increases. Both models predict that when wages are low hired men have greater average ability than hired women because of the first manager caring more about beauty for women. Whether this ability gap disappears depends on the underlying model. When the manager bases decisions on a beauty threshold below which he will not hire, the ability gap disappears because the higher wages make those thresholds increasingly irrelevant. When, instead, the manager trades off beauty and ability when evaluating a prospective employee, the gender ability gap favoring men might persist in the face of rising wages, or it might disappear and then reverse, depending on the distribution of beauty and ability in the population. This suggests that, depending on the nature of the first manager's preferences, gender ability gaps might survive throughout the pay spectrum, and they could go in either direction.

The paper proceeds as follows. Section 2 presents the underlying framework in which the first manager exists in an environment that makes taste-based discrimination possible, and then is succeeded by a new manager who works in an environment where statistical discrimination is possible. Section 3 presents the first model in which a manager has beauty thresholds and does not hire workers whose beauty level falls below that threshold. It shows that when the manager sets different thresholds for the two genders, it lays the basis for subsequent statistical discrimination against the gender

[1] The field experiment evidence in Bagues and Perez-Villadoniga (2012) is consistent with this idea.

with the stricter beauty requirement. Section 4 demonstrates the robustness of this result by presenting an alternative model without beauty thresholds but in which the manager makes different tradeoffs between ability and beauty for the two genders. Unlike the model in Section 3 this one makes particular functional-form assumptions, but like the model in Section 3 it identifies conditions under which the first manager's taste-based hiring provides a foundation for statistical discrimination by subsequent managers. Section 5 discusses testable hypotheses governing how the average ability gap changes as the wage rises, and Section 6 contains some brief concluding remarks.

2. The setting for discrimination

The model requires a setting with two hiring managers, one replacing the other. To fit the paper's goal of demonstrating that taste-based discrimination can leave a legacy of statistical discrimination, we need a setting in which both types of discrimination are possible. Statistical discrimination requires a scenario in which the hiring manager is unsure whether gender matters for productivity or not. Let a denote a worker's ability level, and let the worker's productivity be given by $\theta_g a$, where g denotes gender and θ_M and θ_W are unknown parameters governing how different genders' ability translates into revenue for the firm. To ease interpretation of the results we assume the true values satisfy $\theta_M = \theta_W = 1$ so that there is no basis for gender discrimination, but hiring managers do not necessarily know this.

Because we want any discrimination by the first manager to come from tastes, assume that the first manager knows the true values $\theta_M = \theta_W = 1$, but he possesses tastes over not just productivity (which is the same as ability for him), but also some ascriptive characteristic b which we refer to as beauty. These tastes over productivity and beauty are gender-specific and governed by the utility function $U_g(a,b)$. The manager searches for workers and, consistent with the sequential search literature, when filling an opening he considers in turn a sequence of potential employees and hires the first one who generates utility net of wages above some reservation level, that is,

$$U_g(a,b) - w \ \geq \ u^r \tag{1}$$

Men and women appear in the sequence randomly, and the reservation utility level is assumed to be independent of gender.

Based on these tastes, he hires a workforce in which the average ability of hired women may differ from the average ability of hired men. Let A_g denote the average ability of hired workers of gender g. If the expected average ability of hired men differs from the expected average ability of hired women, this would be a form of taste-based discrimination favoring the gender with higher average ability.

The second hiring manager differs from the first by facing a setting with potential statistical discrimination but not potential taste-based discrimination. Statistical discrimination is only possible if the hiring manager does not know the true values of θ_M and θ_W, and must therefore infer them from the data. Moreover, taste-based discrimination is ruled out if his preferences over workers depend only on perceived productivity and not on beauty. The simplest way to do this is to assume that his preferences are simply $V_g(a) = \theta_g a$, which captures his beliefs about the productivity of a worker with gender g and ability a.

He estimates the relative values of θ_M and θ_W by observing the average productivities of the two subsets of the hired workforce. Because the true values are $\theta_M = \theta_W = 1$, the observed productivity of males hired by the first manager is A_M and the observed productivity of females is A_W. This leads to the inference that:

$$\frac{\theta_W}{\theta_M} = \frac{A_W}{A_M} \tag{2}$$

Like the first manager, the second manager follows a sequential search process when hiring and employs those whose perceived productivity exceed some reservation level:

$$\theta_g \mathrm{a} - w \geq v^r \tag{3}$$

As before, the reservation level is assumed to be independent of gender. It follows that he hires men whenever $a \geq (v^r + w)/\theta_M$ and hires women whenever $a \geq (v^r + w)/\theta_W$. If the first hiring manager's tastes led to a workforce with men having greater average ability than women, the second hiring manager would perceive that $\theta_W \leq \theta_M$, in which case the ability threshold for hiring women would be larger than the threshold for hiring men. This would constitute statistical discrimination on the part of the second hiring manager, and it would occur as a legacy of the first hiring manager's taste-based discrimination. The next two sections detail situations in which the first manager hires a work force with $A_W < A_M$, not because of any outright discrimination against women, but instead because of gender-specific preferences regarding beauty.

Before that two issues merit further discussion. First, allowing the first manager to have taste-based discrimination precludes him from being a profit-maximizer. The worker's productivity is part of the manager's objective function, but so is the worker's beauty, and his willingness to trade one for the other leads to the taste-based discrimination. Even if the firm's principle provided the manager with incentives to maximize profit, the manager's willingness to trade profit for beauty would persist.

Second, the timing of the model allows for two ways for the second manager to avoid statistical discrimination. The first manager could communicate either his hiring standards or the true productivity parameters to the second, in which case the second could correct for them when assessing relative average abilities. Alternatively, the second manager could assess the ability level of each member of the existing workforce and base estimates of θ_g on individual data rather than average data. The legacy of statistical discrimination arises when neither communication between managers nor individual-level assessments occur.

3. Discriminatory minimum acceptable threshold

The purpose of this section is to introduce a model which leads to taste-based discrimination by the first hiring manager. The hiring manager does not value beauty per se, but has beauty thresholds he will not cross. More precisely, consider a manager who observes applicants' gender, ability and beauty, and is unwilling to employ anyone whose beauty falls below a given threshold, no matter how productive they are. We show how different beauty thresholds for male and female applicants lead to gender differences in the characteristics of the hired workforce.

We assume that the manager receives utility when he hires a new worker, and that utility depends on both the ability and the beauty of the hired worker. Consistent with a sequential search setting, the manager has a reservation utility level $u^r > 0$, and he hires workers who surpass the reservation utility but not those who fall below it. The manager's gross (that is, pre-wage) utility from employing a worker with ability (or productivity level) a, beauty level b, and gender g is given by:

$$U_g(a, b) = \begin{cases} u(a, b) & \text{if } b \geq T_g \\ 0 & b < T_g \end{cases}, \tag{4}$$

where $u(a,b) > 0$ whenever $a > 0$ and $b > 0$, $\partial u/\partial a > 0$ for all (a,b), and $\partial u/\partial b > 0$ for all (a,b). This structure has several key components. The gender-specific variable T_g is a beauty threshold below which the worker generates no utility for the manager, making him unwilling to hire. If the applicant's beauty exceeds the threshold, the manager's gross utility is given by $U_g(a,b) = u(a,b)$, and this is gender-independent. Utility strictly increases with worker ability, and also strictly increases with worker beauty. Finally, the fact that θ_g is absent from the utility specification reflects the fact that the manager knows the true values of these parameters, $\theta_M = \theta_W = 1$. We wish to explore the effects of the manager setting different thresholds for the two genders, so we assume that $T_W > T_M$, where W and M denote women and men, respectively.

The manager pays the market wage w to each hired worker, so he hires those who offer net utility $U_g(a,b) - w$ higher than his reservation utility u^r. Job applicants care only about the wage, and they accept the offer as long as $w \geq 0$. Therefore, the condition for an applicant of gender g with ability a and beauty b to be hired is:

$$U_g(a,b) - w \geq u^r$$

which coincides with Expression (1) above.

Job candidates arrive sequentially, with ability and beauty determined by random variables with density functions $f_a(a)$, $f_b(b)$ and distribution functions $F_a(a)$, $F_b(b)$ on intervals [0, A], $[B_1, B_2]$, respectively, where $0 < B_1 < 1 < B_2$. The two random variables are assumed to be distributed independently from each other and also independent of gender. We also assume that the two genders arrive independently and with equal probability, although the second assumption does little but simplify the analysis.

Inserting Eq. (4) into Eq. (1) allows one to find, for each acceptable beauty level b, a critical value of ability above which the manager hires the worker. The critical value $a^*(b)$ solves:

$$u(a^*(b), b) = u^r + w$$

and it applies for any worker of gender g whose beauty level exceeds the threshold T_g.

Differentiating with respect to b yields:

$$\frac{da^*}{db} = -\frac{\partial u/\partial b}{\partial u/da} < 0 \tag{5}$$

Hamermesh and Biddle (1994) found that workers with higher beauty received higher pay. While everyone in our model receives the same wage w, making a beauty-based wage premium impossible, our model does generate two patterns of preferential treatment for those with better looks. First, Expression (5) means that a better-looking person can get hired with lower ability than a worse-looking person. A second version of preferential treatment for the better-looking involves not an individual comparison but a comparison of all those hired. The average ability of a hired worker of gender g and beauty b is given by:

$$H_g(b) = \int_{a^*(b)}^{A} a \frac{f_a(a)}{1 - F_a(a^*(b))} \mathrm{d}a$$

Because all of the components are gender-independent, we have $H_M(b) = H_W(b) = H(b)$ for all b. Differentiating H with respect to b yields:

$$H'(b) = -\frac{a^* f_a(a^*)}{1-F_a(a^*)}\frac{da^*}{db} + \frac{f_a(a^*)}{[1-F_a(a^*)]^2}\frac{da^*}{db}\int_{a^*(b)}^{A} a f_a(a)da = \frac{f_a(a^*)}{1-F_a(a^*)}[H(b)-a^*]\frac{da^*}{db} < 0$$

where the last inequality follows from average ability $H(b)$ being necessarily larger than the lowest hirable ability level $a^*(b)$ along with the result that $da^*/db < 0$. The fact that $H'(b) < 0$ means that better-looking workers are, on average, less able than worse-looking ones. These two results, that a better-looking applicant faces a lower ability requirement for employment and that among hired workers the better-looking are less productive, on average, than the worse looking, are both consistent with the empirical literature supporting preferential treatment based on beauty.

Neither of these results concern gender discrimination, though, because critical ability values for a given level of beauty are gender-independent. Gender discrimination arises because the average ability level for an applicant with beauty level $b \in (T_M, T_W)$ is $H(b)$, and males in this range get hired but females in this range do not. Because $H(b)$ is decreasing, workers in this range have higher average ability than workers for whom $b > T_W$. Consequently, these hired males, as a group, have the firm's highest average productivity. Because women with beauty in this range are excluded, we get the following result.

Proposition 1. *Assume that the manager discriminates based on minimum acceptable beauty with $T_W > T_M$ and that $F_b(T_W) > F_b(T_M)$, so that beauty levels between the two thresholds occur with positive probability. Then the average ability of hired male employees is greater than that of hired female employees.*

Proof. Define

$$A_g = \int_{T_g}^{B_2} H(b)\frac{f_b(b)}{1-F_b(T_g)}db$$

which is the average ability of hired workers of gender g. Then,

$$A_M - A_W = \int_{T_M}^{B_2} H(b)\frac{f_b(b)}{1-F_b(T_M)}db - \int_{T_W}^{B_2} H(b)\frac{f_b(b)}{1-F_b(T_W)}db = [\frac{F_b(T_W)-F_b(T_M)}{1-F_b(T_M)}\int_{T_M}^{T_W} H(b)\frac{f_b(b)}{F_b(T_W)-F_b(T_M)}db$$

$$+\frac{1-F_b(T_W)}{1-F_b(T_M)}\int_{T_W}^{B_2} H(b)\frac{f_b(b)}{1-F_b(T_W)}db] - \int_{T_W}^{B_2} H(b)\frac{f_b(b)}{1-F_b(T_W)}db = \frac{F_b(T_W)-F_b(T_M)}{1-F_b(T_M)}\int_{T_M}^{T_W} H(b)\frac{f_b(b)}{F_b(T_W)-F_b(T_M)}db$$

$$-\frac{F_b(T_W)-F_b(T_M)}{1-F_b(T_M)}\int_{T_W}^{B_2} H(b)\frac{f_b(b)}{1-F_b(T_W)}db = \frac{F_b(T_W)-F_b(T_M)}{1-F_b(T_M)}[\int_{T_M}^{T_W} H(b)\frac{f_b(b)}{F_b(T_W)-F_b(T_M)}db - \int_{T_W}^{B_2} H(b)\frac{f_b(b)}{1-F_b(T_W)}db]$$

The term in brackets in the last line is the difference between the average ability of hired workers with beauty in the range $[T_M, T_W]$ and the average ability of hired workers with beauty in the range $[T_W, B_2]$. Because $H'(b) < 0$, the former is strictly larger than the latter. Combining this with the assumption that $F_b(T_W) > F_b(T_M)$ yields $A_M > A_W$. □

Proposition 1 states that when the manager sets a higher acceptable beauty threshold for one group, in this case women, the average productivity of those women who get hired is lower than the average productivity of hired men. The intuition is straightforward. Because of the high beauty threshold, some women with high productivity but low beauty level will be excluded from employment at this firm. Because they face a lower beauty threshold, the manager might hire some male workers with beauty between T_M and T_W, and these men are likely to have higher ability to make up for their lower beauty. Taken together, these patterns lead to a lower average ability of hired female workers because some low-beauty/high-productivity females would be excluded when their male counterparts would get hired.

As discussed in Section 2, the fact that $A_W < A_M$ means that when the first hiring manager leaves and is replaced by someone who does not know the true values of the parameters θ_M and θ_W, that new hiring manager could try to uncover the values of θ_M and θ_W by observing average productivity of the two genders as in Eq. (2). This leads him to believe that $\theta_W < \theta_M$. He hires new male workers when their ability satisfies $a \geq (v^r + w)/\theta_M$, and hires new female workers when $a \geq (v^r + w)/\theta_W > (v^r + w)/\theta_M$. Consequently, women with ability levels between $(v^r + w)/\theta_M$ and $(v^r + w)/\theta_W$ are not hired while their male counterparts would be. Put differently, if the new manager compares output across genders without correcting for either beauty or ability, he would infer that women tend to be less productive in this task than men. In our setting this relationship is purely correlational, but if he treats it as causal he would have a basis for discrimination against women in his subsequent hiring decisions, even without any preference-based bias of his own. In other words, he would statistically discriminate against women. Thus, the first manager's taste-based discrimination on an ascriptive dimension can lead to a legacy of statistical discrimination at the firm.

4. Discriminatory weight on beauty

In this section we offer an alternative setting which yields the same result—the first manager's taste-based discrimination enables the second manager's statistical discrimination.[2] Consider a new model of taste-based discrimination in which the first manager no longer has minimum beauty thresholds for hiring, but instead cares more about women's looks than men's. More precisely, his marginal rate of substitution between ability and beauty differs between male employees and female employees.

Using the same notation as in Section 3, the utility the manager can obtain from employing a worker with ability a, beauty level b, and gender g is given by:

$$U_M = a^{\alpha} b^{1-\alpha}, \quad U_W = a^{\beta} b^{1-\beta} \tag{6}$$

where $0<\beta<\alpha<1$. The manager has a Cobb-Douglas utility function over ability and beauty, but the coefficients differ between genders with relatively more weight on ability for men and relatively more weight on beauty for women. As before, he offers a job to workers who generate net (after wage) utility that exceeds the reservation utility level u^r. The sufficient condition for an applicant of gender g to be employed is given by:

$$U_g - w \geq u^r \tag{7}$$

where g = M, W.

Inserting Expression (6) into Eq. (7), one can find the critical ability values for hiring men and women with given beauty level b:

$$a_{\mathrm{M}}(b) = \left(\frac{u^r + w}{b^{1-\alpha}}\right)^{\frac{1}{\alpha}} = \left(u^r + w\right)^{\frac{1}{\alpha}} b^{1-\frac{1}{\alpha}} \tag{8a}$$

$$a_{\mathrm{W}}(b) = \left(\frac{u^r + w}{b^{1-\beta}}\right)^{\frac{1}{\beta}} = \left(u^r + w\right)^{\frac{1}{\beta}} b^{1-\frac{1}{\beta}} \tag{8b}$$

To receive a job offer, a male with beauty level b must have ability higher than $a_{\mathrm{M}}(b)$, while a female with the same level of beauty must have ability higher than a_{W}(b). Our analysis proceeds by comparing these two critical values.

As before, job candidates arrive sequentially with ability and beauty determined by random variables with density functions $f_a(a), f_b(b)$ and distribution functions $F_a(a), F_b(b)$ on intervals [0, A] and $[B_1, B_2]$, respectively, where $0<B_1<1<B_2$. The two random variables are assumed to be distributed independently from each other and also independently of gender. We assume that the two genders arrive independently.

Proposition 2. *When $b>u^r+w$, the critical ability value for hiring a male of beauty b is larger than the critical ability value for hiring a female of the same beauty, and the opposite is true when $b<u^r+w$. The critical values are identical when $b=u^r+w$.*

Proof. Comparing critical values for two groups with given beauty level, it can be shown that:

$$a_{\mathrm{M}}(b) - a_{\mathrm{W}}(b) = (u^r + w)^{\frac{1}{\alpha}} b^{1-\frac{1}{\beta}} [b^{\frac{1}{\beta}-\frac{1}{\alpha}} - (u^r + w)^{\frac{1}{\beta}-\frac{1}{\alpha}}]$$

The left-hand side has the same sign as the term in brackets. The maintained assumption that $\alpha>\beta$ implies that $\left((1/\beta)-(1/\alpha)\right) > 0$, which in turn implies that:

$$a_{\mathrm{M}}(b)\begin{cases} > a_{\mathrm{W}}(b) & u^r + w < b < B_2 \\ = a_{\mathrm{W}}(b) \quad \text{when} & b = u^r + w \\ < a_{\mathrm{W}}(b) & B_1 < b < u^r + w \end{cases} \tag{9}$$

□

Proposition 2 identifies two distinct beauty categories which we will call "good-looking" and "plain." Good-looking people have beauty $b \in [u^r + w, B_2)$, while plain people have beauty $b \in (B_1, u^r + w)$. The manager's hiring patterns differ by gender in a way that depends on whether the applicant is good-looking or plain. The manager is willing to sacrifice more ability to get a good-looking female compared to an equally good-looking male, but a plain woman must have more ability than an equally-plain man to secure employment.

As with the threshold model of Section 3, there are two ways in which better-looking individuals can receive preferential treatment. Unlike in that setting, though, here the degree of preferential treatment differs by gender, since the manager is

[2] We think of this as the theory version of a robustness check.

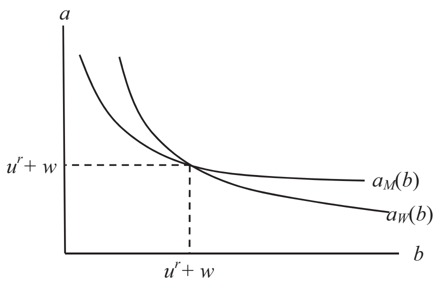

Fig. 1. The minimum required abilities and beauty levels.

more sensitive to female beauty than male beauty. The first type of preferential treatment arises directly from the manager trading off beauty for ability in the hiring decision, so that as beauty increases the required ability level decreases ($da_g(b)/db > 0$). This inequality follows from Eqs. (8a) and (8b) along with the assumptions that α and β are both less than one.

For the second type of preferential treatment, denote, as before,

$$H_g(b) = \int_{a_g(b)}^{A} a \frac{f_a(a)}{1 - F_a(a_g(b))} \mathrm{d}a$$

Thus, $H_g(b)$ is the expected ability of hired individuals of gender g and beauty b. Following the same logic as in Section 3, we find:

$$H_g{}'(b) = \frac{f_a(a_g)}{1 - F_a(a_g)} [H_g(b) - a_g] \frac{\mathrm{d}a_g}{\mathrm{d}b} < 0$$

which means that within each gender the average demanded ability is lower for better-looking people. Thus preferential treatment appears both in terms of the required ability level and the demanded average ability level: both decrease as the applicant's looks improve.

However, unlike the Section 3 case of different minimum beauty thresholds, from Eq. (7) one can see that the critical ability levels differ between genders, and therefore so does the average ability between genders holding beauty constant.

Proposition 3. *As the manager places more weight on women's looks, good-looking men have higher expected productivity than women with the same beauty level, while plain-looking women have higher expected productivity than men with the same beauty level. That is,*

$$H_M(b) \begin{cases} > H_W(b) & & u^r + w < b < B_2 \\ = H_W(b) & \text{when} & b = u^r + w \\ < H_W(b) & & B_1 < b < u^r + w \end{cases}$$

Proof. Write

$$\hat{H}_g(a_g) = \int_{a_g}^{A} a \frac{f_a(a)}{1 - F_a(a_g)} \mathrm{d}a$$

which is the average productivity of gender g workers hired according to the minimum ability standard a_g. Then,

$$\hat{H}_g'(a_g) = \frac{f_a(a_g)}{1 - F_a(a_g)} [\hat{H}_g(a_g) - a_g] > 0$$

where the inequality arises because the $\hat{H}_g(a_g)$ is the average ability of workers when they are hired according to the critical value a_g, and that average value is necessarily larger than the threshold. Thus for a given beauty level b the gender with the higher critical ability level a_g has the higher average productivity $H_g(b)$. For this reason $H_M(b) > H_W(b)$ if and only if $a_M(b) > a_W(b)$, and the result follows. □

Fig. 1 illustrates this result. The two curves show the functions $a_M(b)$ and $a_W(b)$. As is apparent from Eqs. (8a) and (8b), the two curves are decreasing functions of b. They cross when $b = u^r + w$. The employer would hire any male whose combination of ability and beauty lies above $a_M(b)$, and would hire any female whose (a,b) combination lies above $a_W(b)$. For higher levels of beauty, that is, for $b > u^r + w$, the male curve lies above the female one, consistent with Proposition 2, and the ability requirement for a male exceeds that of a female with the same high level of beauty. The opposite pattern holds for plain men and women, with the male curve lying below the female one. When an applicant's ability/beauty combination lies between the two curves the manager is willing to hire them if they are one gender but not if they are the other. For example, for a good-looking applicant falling between the two curves, the manager will hire him if he is male but not if is female.

Because the joint distributions of ability and beauty are identical across the two genders, conditional on being hired a good-looking man has higher expected ability than a woman hired with the same level of beauty. This occurs because lower-ability, good-looking women who lie between the two curves get hired, while all good-looking people lying above the male curve $a_M(b)$ get jobs. Similarly, conditional on being hired a plain-looking woman has higher expected ability than a man hired with the same level of beauty. In both cases, whenever the ability standard differs across the two genders, whichever gender is favored in hiring will tend to have lower expected ability when employed.

In this scenario, as the manager pays more attention to women's beauty, the female curve $a_W(b)$ rotates clockwise, becoming steeper. This results in good-looking women receiving relatively more preferential treatment than equally good-looking men, while plain-looking women are penalized relatively more for their looks than men with the same beauty level. And, as already argued, this increase in attention to women's beauty also leads to relatively more-able set of good-looking men and plain-looking women.

We now turn our attention to the hired workforce as a whole, asking how the gender difference in how the manager trades off beauty and ability affects the average productivity of the men and women he hires. Any differences in the average productivity levels would be caused entirely by hiring practices because there is no underlying difference in the distribution of ability levels across the two genders.

Proposition 4. *If the fraction of good-looking people in the population is high enough, that is, there exists $k \in (0, 1)$ such that if $F_b(u^r + w) < k$, then the pool of hired males has higher average productivity than the pool of hired females.*

Proof. Using the same construction of average hired ability level as in Section 3, the expected ability of male workers is,

$$A_M = \int_{B_1}^{B_2} \int_{a_M(b)}^{A} \frac{af(a,b)}{1-F(a_M(b))} \mathrm{d}a\mathrm{d}b = \int_{B_1}^{B_2} H_M(b) f_b(b) \mathrm{d}b$$

and the expected ability of the female workers is,

$$A_W = \int_{B_1}^{B_2} \int_{a_W(b)}^{A} \frac{af(a,b)}{1-F(a_W(b))} \mathrm{d}a\mathrm{d}b \int_{B_1}^{B_2} H_W(b) f_b(b) \mathrm{d}b.$$

The average ability difference, then, is,

$$A_M - A_W = \int_{B_1}^{B_2} [H_M(b) - H_W(b)] f_b(b) \mathrm{d}b$$

Breaking this into parts,

$$A_M - A_W = F_b(u^r + w) \int_{B_1}^{u^r+w} [H_M(b) - H_W(b)] \frac{f_b(b)}{F_b(u^r + w)} \mathrm{d}b + [1 - F_b(u^r + w)] \int_{u^r+w}^{B_2} [H_M(b) - H_W(b)] \frac{f_b(b)}{1 - F_b(u^r + w)} \mathrm{d}b$$

Define,

$$\Delta_P = \int_{B_1}^{u^r+w} [H_W(b) - H_M(b)] \frac{f_b(b)}{F_b(u^r + w)} \mathrm{d}b$$

to be the average gender ability difference between plain women and plain men. Note that $\Delta_P > 0$ by construction. Similarly, define,

$$\Delta_G = \int_{u^r+w}^{B_2} [H_M(b) - H_W(b)] \frac{f_b(b)}{1 - F_b(u^r + w)} \mathrm{d}b$$

to be the average ability difference between good-looking men and good-looking women, which again is positive by construction. Then,

$$\begin{aligned} A_M - A_W &= F_b(u^r + w) \int_{B_1}^{u^r+w} [H_M(b) - H_W(b)] \frac{f_b(b)}{F_b(u^r + w)} \mathrm{d}b + [1 - F_b(u^r + w)] \int_{u^r+w}^{B_2} [H_M(b) - H_W(b)] \frac{f_b(b)}{1 - F_b(u^r + w)} \mathrm{d}b \\ &= -F_b(u^r + w)\Delta_P + [1 - F_b(u^r + w)]\, \Delta_G = F_b(u^r + w)[-\Delta_P - \Delta_G] + \Delta_G \end{aligned}$$

Set $k = \Delta_G/(\Delta_P + \Delta_G)$. The result follows. □

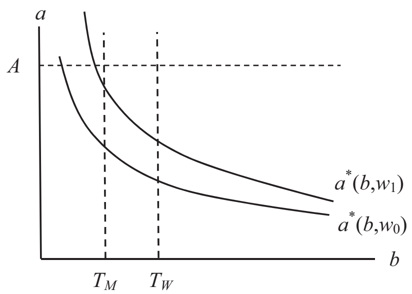

Fig. 2. Minimum required abilities and increased wages, $w_0 < w_1$.

When the manager trades off beauty and productivity, good-looking women tend to be less productive than good-looking men, but plain-looking women have higher productivity, on average, than plain-looking men. Which of these dominates depends on the distribution of beauty relative to the critical level $\hat{b} = u^r + w$. Recalling that men and women have the same underlying beauty distribution, if more men tend to be good-looking then so do more women. Therefore, if enough of the population is good-looking, the average ability loss from good-looking women outweighs the average ability gain from plain-looking ones, leading to the result of Proposition 4.

5. Wages and the gender ability gap: an illustration

This section explores whether and how the gender ability gap changes when the market wage increases. As before, we examine this question using two approaches, one in which the hiring manager has gender-specific beauty thresholds and one in which he has no thresholds but trades off beauty and ability differently for the two genders. Unlike in the previous sections where the two models had the identical qualitative prediction that current taste-based discrimination can generate a basis for future statistical discrimination at the same employer, here the two models have different predictions. In the threshold model the ability gap disappears as the wage increases (holding the thresholds constant), while in the tradeoffs model the ability gap can persist or even reverse.

Begin with the threshold model, but assume as in Section 4 that the hiring manager's preferences take the Cobb-Douglas functional form $U(a, b) = a^{\gamma} b^{1-\gamma}$ where $\gamma \in (0, 1)$. From Eq. (8a) the minimum acceptable ability for a worker of beauty b is:

$$a^*(b) = (u^r + w)^{\frac{1}{\gamma}} b^{1-\frac{1}{\gamma}}$$

so that,

$$\frac{\mathrm{d}a^*(b)}{\mathrm{d}w} = \frac{1}{\gamma}(u^r + w)^{\frac{1}{\gamma}-1} b^{1-\frac{1}{\gamma}} > 0$$

and,

$$\frac{\mathrm{d}^2 a^*(b)}{\mathrm{d}w\mathrm{d}b} = \frac{1}{\gamma}(1 - \frac{1}{\gamma})(u^r + w)^{\frac{1}{\gamma}-1} b^{-\frac{1}{\gamma}} < 0$$

The first of these implies that when the wage rises the manager hires a workforce with greater ability. The second implies that the increases in ability are larger for those with lower beauty levels.

Because of the thresholds, the manager hires men with beauty levels in the range $[T_{\mathrm{M}}, T_{\mathrm{W}}]$ but not women. As can be seen from Fig. 2, these are the workers for whom the critical ability level increases the most, and therefore the beauty levels for which the average worker productivity increases the most. This reasoning suggests that taste-based discrimination leads to even more statistical discrimination in higher-wage industries.

Fig. 2 shows only part of the story, though, because as the wage continues to increase, eventually workers with too little beauty get shut out of the market. Define $b_*(w)$ by the value of b that solves $u(A,b) = w$, where A denotes the maximum ability level. When the wage is w and beauty level is $b_*(w)$, the worker would need the highest possible ability level A to meet the critical ability level for hiring, that is, $a^*(b_*(w)) = A$. Note that $\mathrm{d}b_*/\mathrm{d}w = 1/(\partial u/\partial b) > 0$, and workers with beauty level $b < b_*(w)$ cannot possibly get hired. This leads to two effects as wages rise. On the one hand the increase in wages boosts the relative productivity of the men with beauty levels in the range $[T_{\mathrm{M}}, T_{\mathrm{W}}]$, and this leads to an increase in the relative average productivity of hired males. On the other hand, the increase in wage precludes more workers at the bottom of the beauty scale from being hired. When w is small enough that $b_*(w) \ll T_{\mathrm{M}}$, the first effect outweighs the second, but as $b_*(w)$ approaches T_{W}, the advantaged male workers start dropping out of the acceptable hiring pool. When $b_*(w) \geq T_{\mathrm{W}}$ males no longer have any advantage, and the average productivity of the two genders becomes the same. These trends are

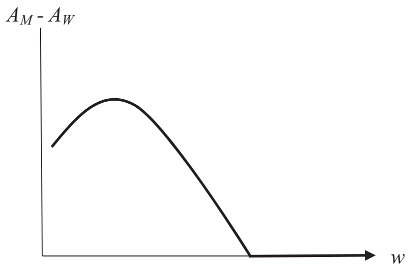

Fig. 3. Wage changes and gender ability gaps in the threshold model.

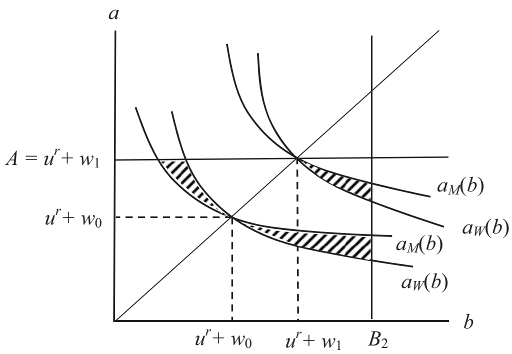

Fig. 4. A persistent gender ability gap with the second model.

summarized in Fig. 3, which shows that the gender ability gap first grows then shrinks and eventually disappears as the market wage increases.[3]

The primary prediction from the threshold model of taste-based discrimination is that as the market wage grows the gender ability gap eventually disappears, consistent with the next proposition.

Proposition 5. *In the threshold model there exists a wage w^* above which the gender ability gap is zero, that is, $A_M - A_W = 0$.*

Proof. Let w^* solve $b_*(w^*) = T_W$. For any $w \geq w^*$ it must be the case that $b_*(w) \geq T_W$ and therefore any hire must have beauty greater than T_W. The thresholds do not matter for such individuals and therefore there are no differences between the expected abilities of men and women. □

This result means that ability gaps are more likely to be observed in low-wage markets than in high-wage ones.[4]

We get a different prediction from the second model with gender-dependent rates of substitution between beauty and ability, though, and Fig. 4 depicts a situation in which the gender ability gap persists as the wage increases. Proposition 2 establishes that for high beauty levels (i.e. $b > u^r + w$) the critical ability level for women is greater than that for men, but for low beauty levels ($b < u^r + w$) the opposite holds. As the wages rise both the $a^*_M(b)$ and $a^*_W(b)$ curves shift outward and become steeper, as before, and the crossing point shifts outward along the 45-degree line. Ability can be no higher than A, while beauty can be no larger than B_2, and in the figure we assume that $B_2 > A$.

Begin with the pair of curves with the lower wage. The shaded region between the curves and to the southeast of the intersection contains ability/beauty combinations for which a woman is hired but a man is not. Thus, hires from this region increase the gender ability gap, and all of the women in this region are considered good-looking according to the endogenous standard of the model. In contrast, the shaded region to the northwest of the intersection contains combinations of ability and beauty for which the man is hired but not the woman, and so hires from this region reduce the gender ability gap. These men are plain-looking according to the endogenous standard. When the wage rises to $w_1 = A - u^r$, the manager can no longer hire plain-looking men from the region to the northwest of the new intersection because no one has ability above A. Because $B_2 > A$ the manager can still hire good-looking women who meet the threshold but have ability/beauty combinations in the shaded region to the southeast of the intersection, and these hires lead to a positive gender ability gap.

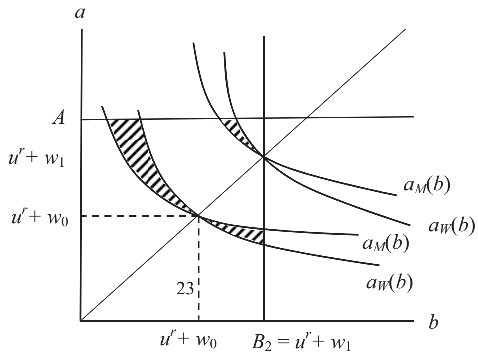

Fig. 5. A persistent negative gender ability gap with the second model.

If $B_2 < A$ the opposite occurs and Fig. 5 shows this case. Once the wage reaches w_1 the demarcation point between plain and good-looking people is B_2, and it is no longer possible to hire good-looking people. Among plain-looking people hired women have greater average productivity than hired men, and the gender ability gap is negative.

Figs. 4 and 5 together show that as the wage grows taste-based discrimination could continue to generate an average ability gap favoring men or it could reverse to generate an average ability gap favoring women. The former occurs if the employer hits the ability constraint before the beauty one, and the latter if the opposite holds. Because managers understand that they should hire based on ability, not the ascriptive characteristic, it stands to reason that the beauty ceiling would be reached first, in which case the average ability gap shrinks and then reverses as the wage increases. If this is the case, then both the threshold and tradeoff models predict a shrinkage in the average ability gap shrinks as the wage increases, but they disagree on whether the gap can reverse.

The persistence of gender ability gaps as wages increase is somewhat counterintuitive. Intuition suggests that as wages grow, discrimination becomes more expensive and therefore the manager would "purchase" less of it. However, the model's use of a reservation utility for hiring keeps the price of discrimination constant. Holding beauty constant, as the wage rises the threshold ability level that makes someone worth hiring increases. Around this threshold ability level, though, the trade-offs between beauty and ability are left unchanged by the wage increase, and so all of the effects of rising wages come from changing the conditional distribution of men and women whose combinations of ability and beauty exceed the hiring threshold.

6. Conclusions

In this paper we construct two types of taste-based discrimination. In the first model a hiring manager sets different group-based minimum acceptable thresholds of an ascriptive characteristic, such as beauty, height, accent, or love of sports, and in the second the manager has different marginal rates of substitution between the characteristic and ability for the two groups. Both models yield the same main result that taste-based discrimination in which the ascriptive characteristic matters more for one group than another leads to a lower average ability of this group. This in turn provides a basis for future statistical discrimination when the original hiring manager is replaced by someone new who observes group differentials in productivity and mistakenly attributes those differentials to group identity and not the previous manager's preferences.

Besides presenting a new channel through which statistical discrimination can arise, the results highlight how seemingly innocuous preference-based hiring can have long-term effects on a company's workforce. Whenever there are two groups of people and the hiring manager cares about some non-productivity-related characteristic more for one group than the other, a legacy of statistical discrimination can emerge. Personnel decisions, then, should either focus on productivity-related characteristics only or else treat all candidates equally regarding ascriptive characteristics.

Appendix A. Construction of Fig. 3.

That variables a, b are distributed uniformly and independently. Set parameters

$$A = 6, B_1 = 1, B_2 = 4, \alpha = \frac{1}{3}$$

[3] Appendix A contains calculations related to the construction of Fig. 3.
[4] This assumes that the thresholds do not change across markets.

Then $f_a(a) = \frac{1}{6}$, $F_a(a) = \frac{a}{6}$, $f_b(b) = \frac{1}{3}$, $F_b(b) = \frac{b-1}{3}$.

$$a*(b) = (u^r + w)^{\frac{1}{\alpha}} b^{1-\frac{1}{\alpha}} = (u^r + w)^3 b^{-2}$$

$$H(b) = \mathrm{d}a \int_{a^*(b)}^{A} a \frac{f_a(a)}{1 - F_a(a^*(b))},$$

$$A_{\mathrm{M}} = \int_{T_{\mathrm{M}}}^{B_2} H(b) \frac{f_b(b)}{1 - F_b(T_{\mathrm{M}})} \mathrm{d}b = \frac{1}{2} A + \frac{\alpha (u^r + w)^{\frac{1}{\alpha}}}{2(2\alpha - 1)} \frac{B_2^{2-\frac{1}{\alpha}} - T_{\mathrm{M}}^{2-\frac{1}{\alpha}}}{B_2 - T_{\mathrm{M}}}$$

$$A_{\mathrm{W}} = \int_{T_{\mathrm{W}}}^{B_2} H(b) \frac{f_b(b)}{1 - F_b(T_{\mathrm{W}})} \mathrm{d}b = \frac{1}{2} A + \frac{\alpha (u^r + w)^{\frac{1}{\alpha}}}{2(2\alpha - 1)} \frac{B_2^{2-\frac{1}{\alpha}} - T_{\mathrm{W}}^{2-\frac{1}{\alpha}}}{B_2 - T_{\mathrm{W}}}$$

a. When $a^*(T_{\mathrm{M}}) \le A$, that is, $w \le 6^{\frac{1}{3}} T_{\mathrm{M}}{}^{\frac{2}{3}} - u^r$.

Let $f(x) = \frac{B_2{}^{2-\frac{1}{\alpha}} - x^{2-\frac{1}{\alpha}}}{B_2 - x}$, $B_1 < x < B_2$, $f'(x) = \frac{B_2{}^{2-\frac{1}{\alpha}} + \frac{x^{-\left(\frac{1}{\alpha}-1\right)}[(1-2\alpha)B_2 - (1-\alpha)x]}{\alpha}}{(B_2 - x)^2}$

If $\alpha \ge \frac{1}{2}$, then $f'(x) \le \frac{B_2{}^{2-\frac{1}{\alpha}} - x^{2-\frac{1}{\alpha}}}{(B_2 - x)^2} < 0$, so $f(T_M) \ge f(T_W)$;

If $0 < \alpha < \frac{1}{2}$, then $f'(x) > \frac{B_2{}^{2-\frac{1}{\alpha}} - x^{2-\frac{1}{\alpha}}}{(B_2 - x)^2} > 0$, so $f(T_M) < f(T_W)$.

Therefore, $\frac{\mathrm{d}(A_{\mathrm{M}} - A_{\mathrm{W}})}{\mathrm{d}w} = \frac{(u^r + w)^{\frac{1}{\alpha} - 1}}{2(2\alpha - 1)} \left(\frac{B_2^{2-\frac{1}{\alpha}} - T_{\mathrm{M}}^{2-\frac{1}{\alpha}}}{B_2 - T_{\mathrm{M}}} - \frac{B_2^{2-\frac{1}{\alpha}} - T_{\mathrm{W}}^{2-\frac{1}{\alpha}}}{B_2 - T_{\mathrm{W}}} \right) > 0.$

b. When $a^*(T_{\mathrm{M}}) > A$, $a^*(T_{\mathrm{W}}) > A$.

Define $\hat{b}$ as $a^*(\hat{b}) = A = 6$, then $T_{\mathrm{M}} < \hat{b} < T_{\mathrm{W}}$, which means that $6^{\frac{1}{3}} T_{\mathrm{M}}{}^{\frac{1}{3}} - u^r < w < 6^{\frac{1}{3}} T_{\mathrm{W}}{}^{\frac{1}{3}} - u^r$.

Then one can get $A_{\mathrm{M}} = \int_{\hat{b}}^{B_2} H(b) \frac{f_b(b)}{1 - F_b(T_{\mathrm{M}})} \mathrm{d}b$,

$$A_{\mathrm{M}} - A_{\mathrm{W}} = -\frac{(u^r + w)^3}{2} \left[\frac{\frac{1}{4} - \sqrt{6}(u^r + w)^{-\frac{3}{2}}}{4 - \frac{1}{\sqrt{6}}(u^r + w)^{\frac{3}{2}}} - \frac{\frac{1}{4} - T_{\mathrm{W}}{}^{-1}}{4 - T_{\mathrm{W}}} \right]$$

$$\frac{\mathrm{d}(A_{\mathrm{M}} - A_{\mathrm{W}})}{\mathrm{d}w} = -\frac{3\left[2(u^r + w)^2 - \sqrt{6(u^r + w)} T_{\mathrm{W}}\right]}{16 T_{\mathrm{W}}}$$

$$\frac{\mathrm{d}(A_{\mathrm{M}} - A_{\mathrm{W}})}{\mathrm{d}w} \begin{cases} \ge 0 & \text{when } w \le \left(\frac{3}{2}\right)^{\frac{1}{3}} T_{\mathrm{W}}{}^{\frac{2}{3}} - u^r \\ < 0 & \text{when } w > \left(\frac{3}{2}\right)^{\frac{1}{3}} T_{\mathrm{W}}{}^{\frac{2}{3}} - u^r \end{cases}$$

If wage is raised so high that $\hat{b} \ge T_{\mathrm{W}}$, according to Proposition 5, the beauty discrimination gap between men and women will disappear.

References

Bagues, M., Perez-Villadoniga, M., 2012. Do recruiters prefer applicants with similar skills? Evidence from a randomized natural experiment. J. Econ. Behav. Organ. 82, 12–20.

Bagues, M., Perez-Villadoniga, M., 2013. Why do I like people like me? J. Econ. Theory 148, 1292–1299.

Becker, G.S., 1957. The Economics of Discrimination. The University of Chicago Press, Chicago & London.

Coate, S., Loury, G.C., 1993. Will affirmative-action policies eliminate negative stereotypes? Am. Econ. Rev. 83, 1220–1240.

Guryan, J., Charles, K.K., 2013. Taste-based or statistical discrimination: the economics of discrimination returns to its roots. Econ. J. 123, 417–432.

Hamermesh, D.S., 2011. Beauty Pays: Why Attractive People are More Successful. Princeton University Press.
Hamermesh, D.S., Biddle, J.E., 1994. Beauty and the labor market. Am. Econ. Rev. 84, 1174–1194.
Hersch, J., 2008. Skin color, immigrant wages, and discrimination. In: Racism in the 21st Century. Springer, New York, pp. 77–90.
Johnston, D.W., 2010. Physical appearance and wages: do blondes have more fun? Econ. Lett. 108, 10–12.
Kinzler, K.D., DeJesus, J.M., 2013. Northern = smart and Southern = nice: the development of accent attitudes in the United States. Q. J. Exp. Psychol. 66, 1146–1158.
Lang, K., Lehmann, J.K., 2012. Racial discrimination in the labor market: theory and empirics. J. Econ. Lit. Am. Econ. Assoc. 50, 959–1006.
Liu, X., Sierminska, E., 2014. Evaluating the Effect of Beauty on Labor Market Outcomes: A Review of the Literature. IZA DP No. 8526, October.
Morgan, J., Vardy, F., 2009. Diversity in the workplace. Am. Econ. Rev. 99, 472–485.
Persico, N., Postlewaite, A., Silverman, D., 2004. The Effect of Adolescent Experience on Labor Market Outcomes: The Case of Height (No. w10522). National Bureau of Economic Research.
Phelps, E.S., 1972. The statistical theory of racism and sexism. Am. Econ. Rev. 62, 659–661.