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

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.

© 2016 Elsevier B.V. All rights reserved.

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

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.