Can Statistical Discrimination Explain Inequality?

Christian Hunkler, University of Mannheim

- first draft -

Abstract: Statistical discrimination starts from the notion of employers' incomplete information about the real productivity of applicants, even some time after hiring. The paper focuses on the disputed question whether Phelp's measurement model of statistical discrimination can explain inequality in hiring, i.e. group discrimination. The central argument is that employers perceive differences in the reliability of productivity signals between groups. They may trust productivity signals less (e.g. test scores), if these signals come from a specific group of applicants. The *theoretical analysis* finds the model not being capable to explain inequality in wages. However, when considering allocation decisions (the hiring of workers) the mechanism may indeed result in group discrimination, i.e. inequality. Contrary to intuition the direction of discrimination depends on the relation of workers seeking a job to the number of open positions. Using *simulations*, I show that the group whose signals are trusted less is discriminated in very competitive labor markets, whereas, under the conditions of less competition, the model even predicts discrimination against those workers whose signals are trusted more. Considering the access to qualified positions, however, the model always results in discrimination against the group whose signals are trusted less, regardless of the level of competition. The paper concludes with some general considerations on discrimination theory.

Contents

1. Introduction: Statistical Discrimination

Statistical discrimination starts from the notion of employers' incomplete information about the real productivity of applicants, even some time after hiring (Aigner and Cain 1977; Arrow 1973; Arrow 1998; Oaxaca 2001; Spence 1973). Since employers face an investment decision under uncertainty, they base their decisions on the interpretation of easy accessible characteristics. Spence (1973: 357) introduced two kinds of characteristics: *indices*, being traits that are observable and not changeable such as gender or race; and *signals* referring to observable and changeable traits, the most prominent being educational degrees. Inferences from indices and signals about applicants' actual productivity are possible because of employers' former experiences with other employees. Hence they possess knowledge about the general distribution of productivity. For example, employers may hold the "statistical belief" that applicants who possess university degrees are more productive in a certain job than other applicants who only hold high school degrees. There are three different models of statistical discrimination (Correll and Benard 2006; England 1992; Kalter 2003). The *first* elaborates on differences in statistical beliefs on group average productivity; the *second* on differences in variances of (same) group productivity. I exclusively focus on the *third*, the so called measurement model of statistical discrimination, as I agree with previous reasoning that the first two can only explain group inequality if additional and questionable assumptions are introduced (England 1992; Kalter 2003; Kalter 2006).

The core mechanism proposed in the measurement model of statistical discrimination is a differential in employers' perception of the reliability of a given productivity signal. The disputed question is, whether or not inequality arises, when employers trust the same productivity signal from one group less. To clarify, inequality is meant in the sense of group discrimination, i.e. when of two groups who provide "essentially identical" labor services [1], one group on average earns less or its members have a systematically lower chance of being hired. I am especially concerned with ethnic inequalities, however, the following analysis should be generalizable to any form of group inequality/discrimination. In the field of ethnic inequality, usually ethnic minorities' labor market disadvantages are attributed to a lack of productive resources on various dimensions (e.g. Heath and Cheung 2007; Kalter 2006). In addition, there is a debate on whether ethnic

[1] That is no pre-market differences in average productivity or skills or availability.

minorities are the target of discriminative hiring behavior by employers (Seibert and Solga 2005; Kalter 2006). This paper tries to clarify whether statistical discrimination can explain group inequality, with a special focus on the explanation of ethnic minorities disadvantages. For the purpose of illustration, I assume a minority group B's signals to be perceived as less reliable in comparison to some majority group W's signals.[2]

2. The Measurement Model of statistical discrimination

The measurement model of statistical discrimination was introduced by Phelps (1972), usually the more accessible description by Aigner and Cain (1977) is cited. It assumes employers to base wage or hiring decisions on visible signals (y), as long as they believe these signals are reliable "test scores" of the true productivity q they are interested in. The underlying reasoning is that signals, e.g. aptitude test scores or grade point averages, are a function of the true productivity q and a normally distributed error term u as shown in formula 1:

$$y = q + u \tag{1}$$

To the extent that employers believe a signal to be a less reliable test score, i.e. the error term in the regression like relation implied in formula 1 is larger, they rather go back to their statistical beliefs about average group productivity α. Therefore, the central parameter in the measurement model is the reliability γ of visible signals, which may be group specific (indexed by g). It expresses the extent to which employers use signals or group averages to form an opinion about expected productivity, as implied in formula 2.

$$E(q|y) = (1 - \gamma_g) * \alpha_g + \gamma_g * y \tag{2}$$

[2]There are good reasons to expect majority employers to trust signals of minority groups less. The general argument of statistical discrimination is that past experiences are the basis of statistical beliefs. The measurement model assumes non-distorted perceptions of average productivity and of signals; an assumption questionable in itself, considering everything we know about distortions in cognitive perceptions (see Hunkler 2008, especially chapter 2.3.2). However, this would suggest "error-discrimination" (England 1989) that should erode by arbitrage in the long run. Nonetheless, it is likely that the reliability of minority members' signals is perceived lower, as employers by definition have less experience with *minority* candidates. As long as this minority status remains, the hereon based lower reliability perception by employers may not erode by arbitrage, as easily as erroneous beliefs about average productivity would.

Formula 2 describes employers target variable, they are interested only in the expected value of true productivity q given the observable test score y, i.e. $E(q|y)$. The formula is derived as a reverse regression based on formula 2 assuming q and u are joint-normally distributed and not correlated. To make inferences about $E(q|y)$ for a worker or applicant from a group, where the signal for some reason is found to be less reliable, employers need to have some statistical beliefs about the reliability of the signal denoted by γ, and about the average productivity of groups, denoted as α. Rational employers should rely more on the test score or signal, when the reliability of this signal is higher, as signals are a more individual information than group averages can be.

The crucial and disputed question is whether this formulation of statistical discrimination implied in formula 2 can actually explain group inequality. More precisely, will a minority group for whom employers assume the same productivity signal as less reliable in fact be on average disadvantaged? Cain (1986: 724) and others (e.g. England 1992; Kalter 2003) agree that individual miss-allocations are unavoidable due to the error term u, however, that group inequality does not follow from the model. I agree that this is true when applying the model to wage determination as Cain did (1986: 723). However, applied to hiring processes, I come to a different conclusion.

2.1. Statistical Discrimination in Wages

To follow Cain's argument on wage determination (1986: 722f.) let us assume a minority and a majority group of workers who on average have the same productivity and employers who have valid statistical beliefs about this average productivity. Using the classical subscripts W for the majority group and B for the minority one obtains $\alpha = \alpha_W = \alpha_B$. This assumption is crucial, as a differential in real average productivity would imply pre-marked differences and hence not market discrimination. Pre-market difference may exist, e.g. unequal chances in acquiring the necessary qualifications in the schooling system, but if employers select on the basis of real differences in productivity this cannot be regarded as discrimination. Assuming valid perceptions by employers is central as well, as otherwise employers with erroneous beliefs should be driven out of markets by arbitrage in the long run (England and Lewin 1989: 242f.). The only difference between majority and minority workers showing the same signal y then is a lower reliability for the latter ones, i.e. $\gamma_W > \gamma_B$. If employers use expected productivity, as implied in formula 2, to determine wages for Ws and Bs Cain derives formula 3:

$$wage_W - wage_B = (y - \alpha)(\gamma_W - \gamma_B) \tag{3}$$

It is obvious from formula 3 that there is no group inequality or group discrimination in wage determination. For a given signal y that is above the mean productivity α majority members will earn a wage premium, however, for a given signal below the mean productivity α minority members will earn a higher wage. On average these will cancel each other out. This can also be seen in formula 2: If mean perceived productivity α was the same, and signals y are non-distorted measures of productivity, there is no effect on average if employers use the signal for one (majority) group and their statistical beliefs for the other (minority) group.

2.2. Statistical Discrimination in Hiring

Let us now turn to hiring processes, when employers decide based on the expected productivity as implied in formula 2 whom of the applicants to hire for an open position in their company. All the straight-forward assumptions of the above wage determination analysis remain unchanged. The only addition is that employers have a choice between different candidates, of whom one is hired[3]. Matters are more complex in a hiring situation for two reasons: First, there is variety in the quality of positions, which usually coincides with differences in productive skills to qualify for these positions. Hence, there may be two levels of inequality: first, who gets a position at all, and second, who gets a qualified position. Second, there can be varying excess demand for positions, i.e. the relation between workers seeking a job versus job-openings may differ [4]. I start with a simplistic example to illustrate how these additional conditions affect whether or not the measurement model of statistical discrimination causes inequality and what kind of inequality. Based on this example, I distinguish two situations and derive hypotheses on the

[3]The analysis can easily be adapted to employers who hire more than one worker at a time.

[4]Throughout this paper I assume that there are always more candidates seeking a job than job-openings available. If there are as many job-openings as candidates seeking a job or even more, this simple model implies that everyone will be hired eventually and no inequality can be present by definition.

resulting patterns of inequality.

2.2.1. A Simple Example

Assume a simple labor market consisting of four candidates, two of them belonging to a group W and the others to a group B. In both groups there is a worker with a signal above the mean (workers 1 and 3) and one below the mean (workers 2 and 4). Hence, there is again no pre-market inequality implied, for both groups the mean signal α is 50. Table 1 shows these four workers, column 3 contains their original signals, column 4 displays the modified signals perceived by statistically discriminating employers. Employers perceive signals from B candidates as less reliable ($\gamma = 0.4$) than signals for W candidates ($\gamma = 0.4$).

Table 1: A simple four worker example

(1) Group	(2) Worker	(3) Signal Y	(4) Signal Y_{SD}^1	(5) α	(6) γ
W	1	60	58	50	0.8
	2	40	42		
B	3	60	54	50	0.4
	4	40	46		

[1] Signal Y_{SD} is computed using formula 2; γ and α are set according to values given in in columns 5 and 6.

Without differences in necessary skill levels the level and direction of inequality solely depends on the excess demand for labor. If employers want to fill exactly two positions, there is no inequality. Worker 1 and 3, hence one W and one B worker will be hired. However, if the relation of open positions to workers is smaller than 1:2 there is an advantage for the W worker with the above mean signal. This is due to the remaining advantage of majority members with signals above the mean α (for W worker 1). He is evaluated rather on the basis of his signal and hence will outdo the comparable minority applicant (worker 3) who is evaluated rather on the basis of the mean productivity. In this case there will be discrimination against the group of workers, whose signals are evaluated as less reliable. If the relation of open positions to workers is larger than 1:2 this is no longer the case. Indeed, under this condition, workers 1, 3, and 4 will be hired. That is two B workers and one W worker, hence there is discrimination against the group of workers, whose signals are evaluated as more reliable.[5] At first glance this is counterintuitive.

[5]In the trivial case of a perfect labor market, where the relation of open positions to candidates is 1:1 there cannot be inequality, because every worker is hired.

However, it is exactly the same compensating effect of the advantage of those minority workers with signals below the average productivity we saw in the wage analysis. The crucial difference in a hiring situation is that this compensating effect does not automatically emerge. Rather it solely depends on the share of available workers hired.

In a next step we broaden the analysis to the existence of *different job qualities and varying skill levels associated with job quality.* For example a good job could require a skill level of at least a signal of 55. It is obvious that in the above 4-worker example only W-worker 1 will be hired. Regardless of the relation of open positions to job-seekers, "Ws" now have an advantage to occupy the good positions, as the "majority" advantage to be evaluated more according to the actual signal plays out, whereas the "minority" advantage has no effect for these positions. To generalize, if good positions require a skill level above the average skill (or signal) level of α discriminatory hiring against the groups whose signals are trusted less results. The level of inequality should rise the higher the required skill level for qualified positions is. For jobs requiring a lower skill level, again the relation of *remaining* W and B workers is relevant; that is, the above described relations of applicants to open positions will determine the level of inequality considering who gets a job at all.

2.2.2. Hypotheses

In a singular hiring situation for a single open position, the boundary condition for the measurement model of statistical discrimination to explain inequality in hiring against applicants whose signals are trusted less, is the presence of at least one worker in the applicant pool with a signal above the mean productivity whose signal is more trustworthy. Considering more than one open position to fill and varying competition about these positions this simple deduction does not hold anymore. Using the above example, I derive general hypotheses about the conditions under which the measurement of statistical discrimination will cause what kind of inequality/group discrimination:

> *Hypothesis 1:* In case of no differences in necessary skill levels,
>
> *(a)* if the relation of open positions to workers is smaller than 1:2 , we will observe discrimination against the group whose signals are trusted less ("high competition"

condition).

(b) if the relation of open positions to workers is equal to 1:2 we observe no discrimination.

(c) if the relation of open positions to workers is larger than 1:2 (but smaller than 1:1 obviously) we will observe discrimination against the group whose signals are trusted more ("low competition" condition).

Hypothesis 2: In case of differences in job quality and varying skill levels associated with job quality, regardless of the relation of open positions to workers, the group whose signals are trusted more will have an advantage to occupy good positions. This advantage increases, the higher the threshold of skills required for a qualified job is.

2.3. Theoretical Conclusion

In sum, the measurement model of statistical discrimination can explain inequality in employment or in quality of job position or job-status, depending on the structural conditions in the respective labor market segment. Only in the very specific case of a relation of applicants to open positions of 1:2, the measurement model does not predict inequality. The mechanism thereby does neither rely on pre-market differences in average productivity, nor on a disadvantage in signaling power of minority members. It solely explains inequality by the assumption of employers putting different levels of trust in the signals of applicants from different groups. Contrary to expectations the mechanism does not only predict discrimination against the group whose signals are trusted less. Concerning the case of no differences in necessary skill levels and low competition, the mechanism would even predict lower employment shares for the group whose signals are trusted more.

3. Simulating Statistical Discrimination

For lack of sufficient empirical data on actual hiring processes including employers decision parameters and workers exact characteristics, I use simulations to show that the above outlined measurement model of statistical discrimination can explain group discrimination in hiring. This method allows to control employers and workers parameters carefully and thereby to reject alternative explanations. Hence, I can attribute found group discrimination, i.e. inequality, clearly to the measurement model's mechanism.

3.1. The Simulation Model

The statistical package Stata is used to simulate sets of workers ($L_{i \cdots n}$) and employers ($E_{j \cdots m}$). Using a standard data matrix, the first columns of the matrix represent the workers' traits. Each employers' decision parameters are written into the first rows of a new column for each employer (these are group-specific α and γ terms). This creates a matrix as shown in table 2 for each single run of the simulation. The cells, where workers in lines intersect with employers in columns, are used to calculate different hiring functions, as described below. An "x" represents an application; e.g. workers L_1 and L_i apply to employer E_1. Hence, employer E_1 in the example can only choose to hire one of the two.

Table 2: Example Matrix of Data Setup

(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
				E_1	E_2	...	E_j	...	E_m
			α_W	50	50	...	50	...	50
			α_B	50	50	...	50	...	50
			γ_W	0.8	0.8	...	0.8	...	0.8
			γ_B	0.6	0.6	...	0.6	...	0.6
	Signal	Group							
L_1	50	W		x			x		
L_2	45	B			x				
...	...	...							
L_i	55	B		x					
...	...	...							x
L_n	50	W					x		

3.1.1. Actors: Workers & Employers

Each worker is assigned a productivity signal y drawn from a standard normal (Gaussian) distribution with a mean of 50 and a standard deviation of 10. Then a random variate function, that allows for different shares of W and B workers, divides the n workers into two groups, called W and B. Assigning the signal first and then allotting each worker to one of the groups should result in almost equally distributed parameters for the two groups.

For each employer an additional column is added to the matrix (column 4 in table 2 is merely used for descriptive purposes, column 5 holds the first employer). The first five rows of the matrix are reserved to store the employers' identification numbers ($E_1 ... E_n$) and their individual decision parameters. As two groups of workers are used, employers may have two different perceptions of mean productivity α (α_W and α_B). However, to make sure that there is no difference between α_W and α_B, both parameters are set equal to the average productivity of all workers (for each employer). For the central parameter γ as well, employers can have different perceptions of how reliable the same productivity signal is from each group. In the example matrix in table 2 each employer has the same perception of mean productivity for both groups, but signals from the W group of workers are perceived as more reliable (0.8) than signals from the B group (0.6).

3.1.2. The Application Process

The application process is designed to resemble the structural features of real world labor markets. The workers only apply to a specified number of employers (in the example matrix in table 2 worker L_1 applies to employers E_1 and E_j). Again a random variate function is used to draw a set of employer numbers for each worker.[6] Employers can only choose from workers who "selected" to apply to them. In the example matrix employer E_1 can only choose to hire either worker L_1 or worker L_i. To keep it simple, each employer has only one open position to fill. In order to compare statistical discrimination to a "fair" hiring routine this step is performed with two selection functions. In the *highest signal hiring function* employers simply choose based on the raw signals and select the applicant with the highest signal; for the *statistical discrimination hiring function* the signals are modified first, using the employers' decision parameters according to formula 2. When an employer hires a worker all other applications of this worker to other employers are deleted, as he is not available on the market anymore. This reduces the set of applicants to the next employers who hire. The order of employer decisions is simply determined by their identification numbers. However, as employers are randomly set up and applications are randomly assigned, the ordering has no impact on the central mechanisms.

[6]This may result in a worker getting assigned two or more times to the same employer. In that case only one application of this worker to the employer is considered.

3.1.3. Results and Robustness Issues

The central outcome of each simulation run is computed as share of W and B workers employed under each of the two selection functions. The finite amount of workers and employers results in small differences regarding the distribution of parameters, e.g. workers signals. Even if the groups are distributed equally (50% Ws and 50% Bs) and randomly chosen from one distribution of signals, standard deviations for one group can get larger and/or the mean signal might be slightly different. To control for this, each combination of parameters (each run of the simulation) is repeated 20 times. To better distinguish outcomes of single simulation runs in figures or plots a small random term on the x- and y-dimension is added.

3.1.4. A Basic Check of the Simulation Model

To show the simulation model works and that simulated inequalities can be attributed to the measurement model of statistical discrimination, in this first step, I use a simple parameter set-up. Then I compare the two selection functions, the *highest signal hiring function* and the *statistical discrimination hiring function.* Panel A in table 3 gives an overview of the parameter settings. 500 workers compete for 100 employers, i.e. 100 open positions to fill. The two groups are equally large, 50% of the workers are assigned to group B and 50% to group W; each worker in both groups applies for 10 positions. For W-workers the employers assume a high reliability of the signal, γ_W is set to 0.8 throughout all runs. The only parameter that varies is γ_B; it is set in 0.05 steps within the interval [0.2; 0.8], creating 13 different γ_B values. Multiplied by 20 repetitions for each parameter setting, the simulation is repeated 260 times. If the above theoretical analysis is correct, hypothesis 1a should hold as the relation of open positions to workers is smaller than 1:2. Under this condition a smaller γ for one group should result in group discrimination against this group. That is their chances of getting a job are lower. Therefore, I expect increasing inequality the smaller γ_B gets in relation to γ_W, in direction of discrimination against B-workers.

Panel B of table 3 shows the simulated signals [7]. Across all 260 single runs, on average the W-workers mean signals and the standard deviation of their signals is somewhat higher. However,

[7] Basis for Panel B in table 3 are the means and standard deviations of the assigned signals of each worker extracted from each single simulation run. Therefore, table 3 e.g. reports standard deviations of the distribution of these original standard deviations.

Table 3: Summary statistics

	Mean	Std. Dev.	Min.	Max.
A: Parameter Settings				
E - Number of Employers	100.0	0	100.0	100.0
γ_W	0.8	0.000	0.8	0.8
γ_B	0.5	0.187	0.2	0.8
L - Number of Workers	500.0	0	500.0	500.0
Share of B Workers	50.0	0	50.0	50.0
Number of Applications Ws	10.0	0	10.0	10.0
Number of Applications Bs	10.0	0	10.0	10.0
B: Simulated Values				
All Workers: Mean Signal	49.991	0.430	48.704	51.478
W Workers: Mean Signal	50.014	0.641	48.445	51.994
B Workers: Mean Signal	49.967	0.641	48.212	51.774
All Workers: Std. Dev. Signal	9.986	0.318	8.963	10.771
W Workers: Std. Dev. Signal	9.984	0.441	8.923	11.026
B Workers: Std. Dev. Signal	9.976	0.421	8.504	11.201
Number of Runs		260		

these very small differences (0.04, respectively 0.01) on a signal range from about 3 to 94 are negligible. Regressing the share of employed B-workers under the statistical discrimination hiring function on the delta in γ (γ_W-γ_B) results in a β-coefficient of $-.23$. Adding controls for mean and standard deviation of the Bs and Ws signals does significantly improve the model-fit. However, the largest effect is only $\beta = -.02$ for the mean B-workers' signals. The effect of the delta in $\gamma_W - \gamma_B$ completely remains (see Appendix A)[8].

Figure 1 shows the aggregated output of the 260 simulation runs. *Panel A* sums up all runs when the *highest signal hiring function* is used. The x-axis is merely used to spread out the runs to make them more distinguishable, but as γ is not used in this hiring function it has no substantive meaning. It becomes obvious from panel A that the random assignment of signals for the two groups was successful. Each groups' employment rate is 20%, which resembles the relation of all open positions (100) to the total number of workers (500) in these simulation runs. There is some "noise" in that only on average (cf. the median bands) the employment share is 0.2, but in single runs and for single groups the share of employed can be as low as 14% or as

[8]The opposite regression on the share of employed W-workers comes to the same conclusion, see Appendix A

high as 26%. This goes back to the finite number of workers and the therefore existing small differences in the distribution of signals, and/or the additional random procedures implied in application of workers to employers. However, the "noise" is reasonably small[9] and the use of 20 repetitions is sufficient to have no considerable disturbances in the median bands.

Figure 1: Simulation Runs of Basic Model

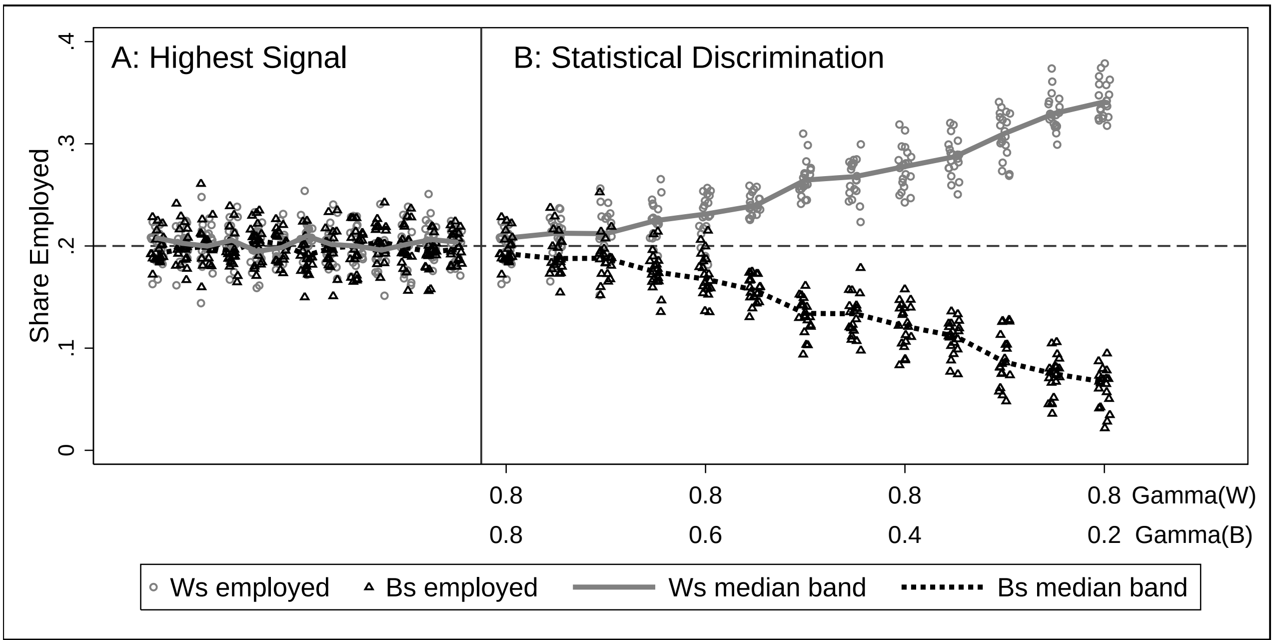

In Panel B of figure 1 the very same signals are the basis for employers' decisions, only now they are modified according to the measurement model's assumption (cf. formula 2). To safely attribute possible effects to the γ-mechanism, instead of using group specific α-terms, α was fixed for both groups to all workers mean productivity within each run. For W-applicants γ_W is set to 0.8 in all runs, whereas γ_B varies from 0.8 to 0.2. The runs are plotted against the different γ values on the x-axis. In the left segment of Panel B, where both groups are assigned the same γ, we should not see any difference. The larger the difference between the perceived reliability of signals gets, i.e. the more B-workers are evaluated according to the mean productivity instead of individual signals, the more W-workers should benefit from their advantage when they have individual signals above the mean. That is exactly what we see in Panel B: with a rising delta

[9] The difference in $\gamma_W - \gamma_B$ explains 83% of the variance in employment of B-Workers, and as well 83% in the employment of W-Workers under the statistical discrimination hiring function. So the randomness implied in the simulation procedure and the non-perfect assignment of productivity signals accounts only for less than 20% of the variance observed (cf. appendix A).

between group specific γ-terms, the share of employed W-workers rises to more than 30%, and the share of B-workers drops below 10%.

3.2. Results: Simulation of Access to Positions

Hypothesis 1 predicted the measurement model of statistical discrimination to result in inequality depending on the relation of open positions to workers. In a "high competition" condition, i.e. the relation of open positions to applicants is smaller than 1:2, I expect discrimination against the group whose signals are trusted less (Bs in the following). In a labor market, where the relation of open positions to workers is equal to 1:2, we will observe no discrimination. Finally in a "low competition" condition, i.e. the relation of open positions to workers is larger than 1:2, discrimination against the group whose signals are trusted more (Ws) is expected . To test this hypothesis I broaden the above used parameter set-up. Instead of using a fixed number of open positions, the amount of available jobs ranges between 50 and 450; the number of workers remains constant (500). To check if the effects are robust concerning the relative size of a minority group whose signals are perceived as less reliable, I additionally vary the relative size of group B. In one set of runs 10% of all workers are Bs, in a second set 30%.

[figure 2 about here]

Figure 2 summarizes the different parameter settings. The first row of figures shows a high competition condition (Hypothesis 1a), the middle row shows a relation of 1:2 (Hypothesis 1b), and the bottom row shows a low competition condition (Hypothesis 1c). In the left panel the share of Bs is 10%, in the right panel 30%. The derived hypotheses all hold. In a high competition condition (first row), there is clear discrimination against the group whose signals are perceived less reliable (against Bs). The lower the reliability parameter γ_B gets, while γ_W stays stable at 0.8, the lower are the simulated employment rates for Bs. At the same time, the employment rates of Ws slightly increase. The increase is small and somewhat difficult to see, due to the W group being very large in comparison to the B group, whose members "lose their jobs" to W workers. This basic finding is robust across different minority sizes, however, it seems that small minorities suffer more under this condition. In the left panel where the minority consists of only

10% of the workers, their medium employment rate is always below the employment rate in the right panel, where the relative minority size is larger (30%). Clearly, for the relation of 1:2 of open positions to workers there is no discrimination visible, regardless of the minority's relative size. Thus, hypothesis 1b perfectly holds. Hypothesis 1c predicted a reversal of the direction of discrimination under the condition of low competition. In the bottom row of figure 2 exactly this can be observed. When the relation of applicants to open positions is larger than 1:2, the group whose signals are trusted more, the Ws, fare worse. Even if the Ws with signals above the average productivity are hired for the first 25% of open positions, the group whose signals are trusted less has an advantage concerning the filling of the following 50% of jobs. Therefore, under the condition of low competition, the prediction of the model reverses, i.e. we observe discrimination against Ws, the group whose signals are trusted more.

3.3. Results: Simulation of Access to Qualified Positions

For access to qualified positions, hypothesis 2 stated that the measurement model results in discrimination against the group whose signals are trusted less, the higher the threshold of skill required for a qualified position is.[10] To test for this hypothesis the above used application procedure was slightly changed. Employers only hire an applicant for a qualified position if the perceived productivity of the respective worker is above a certain threshold.

[figure 3 about here]

Figure 3 shows the simulations for access to qualified positions. For the left panel the threshold of minimum required skills is set to 53; in the right panel it is set to 56[11]. The first row presents a relation of open positions to workers of 50 : 500; the middle row of 250 : 500 and the bottom row of 450 : 500. Throughout the simulated parameter space, the group whose signals are trusted more shows a significantly higher share of being employed in qualified positions (all $t > 2.81$, all $p < 0.01$; $df = 199$). Only for the simulation runs when $\gamma_W = \gamma_B = 0.8$ the employment share is

[10] I restrict the analyses to the usual case of qualified positions requiring skill levels above average productivity.

[11] Using higher thresholds results in similar effects, however, the share of Bs employed drops to zero for larger γ_B values. Zero shares create some kind of bottom ceiling, which makes interpretation unnecessary difficult. Therefore, I use comparatively low thresholds to show the basic mechanism.

on average almost equal (the difference across all runs is -0.003, pointing to a small advantage for Ws; $t = 1.07$,$p > 0.2$, $df = 199$). Therefore, hypothesis 2 cannot be rejected; under all conditions simulated the group whose signals are trusted less is discriminated concerning the access to qualified positions. Figure 3 additionally shows some interesting patterns. In the first row of simulations Ws seem to "take over" Bs qualified positions, the less reliable employers perceive B's signals (the grey Ws median band line increases here, whereas in the other 4 simulations it remains on a stable level). This is due to the small number of only 50 qualified positions in relation to 500 applicants. There are simply enough Ws who show a signal above the thresholds (53 respectively 56), when only 50 positions are to be filled. If a higher share of qualified positions is open (middle and bottom row of simulations) already all Ws whose signal is above the respective thresholds are employed, hence they cannot "win over" the positions B's used to occupy. This is especially obvious when comparing the middle to the bottom row of simulations in figure 3: the number of open positions increases by 200, however the share of Ws employed remains at about 35% for a threshold of 53, respectively 23% for a threshold of 56 (the absolute number of Ws does not change).

The second part of hypothesis 2 predicts more discrimination against Bs the higher the threshold of skills required for a qualified job is. From the comparison of the left panel of figure 3, where the threshold is 53 and the right panel, where the threshold is 56, this prediction is not easy to see. The reason is that with a higher threshold the overall employment level is lower and the figures cannot be directly compared. In addition in some simulations the share of Bs employed in qualified positions even hits 0% at some point. Odds ratios are suitable to show the effect of a higher threshold independent of the absolute shares of hired workers. Suppose S_W and S_B are the shares of Ws respectively Bs who got a qualified position (the plotted values in figure 3), then the odds ratio in favor of Ws is defined as $(S_W/(1-S_W))/(S_B/(1-S_B))$. An odds ratio > 1 indicates an advantage for Ws; an odds ratio < 1 indicates an advantage for Bs. Using odds ratios allows an interpretation of the relative increases in the disadvantages, regardless of the differing absolute levels of the share of employment. However, the odds ratio is only defined as long as no share drops to zero. For the simulation runs in figure 3 for all $\gamma_B > 0.45$ the shares do not drop to zero. Therefore, table 4 shows the average odds ratios for these parameter combinations only. With few exceptions (these being all in the simulations with the extreme combination of 50 positions to 500 workers), the advantage for Ws is larger for the higher threshold of 56. Hence, the second part of hypothesis 2 also holds: the higher the threshold the more will a group, whose signals are trusted less, suffer from statistical discrimination.

Table 4: Simulation of Access to Qualified Positions

		Number of Employers and Threshold														
		—50—			—150—			—250—			—350—			—450—		
γ_W	γ_B	53		56	53		56	53		56	53		56	53		56
.8	.5	5.17	$<$	6.58	1.51	$<$	2.41	1.44	$<$	2.56	1.82	$<$	2.20	1.60	$<$	2.31
.8	.55	4.64	$>$	3.09	1.48	$<$	2.08	1.22	$<$	2.00	1.43	$<$	2.15	1.29	$<$	1.89
.8	.6	2.60	$<$	3.05	1.43	$<$	1.75	1.32	$<$	1.56	1.18	$<$	1.93	1.34	$<$	1.63
.8	.65	2.05	$>$	1.66	1.25	$<$	1.67	1.26	$<$	1.47	1.25	$<$	1.39	1.23	$<$	1.51
.8	.7	1.39	$<$	1.78	1.30	$>$	1.23	1.03	$<$	1.31	1.14	$<$	1.30	1.09	$<$	1.27
.8	.75	1.37	$=$	1.37	1.06	$<$	1.18	1.02	$<$	1.11	1.03	$<$	1.07	1.06	$<$	1.23
.8	.8	1.20	$>$	1.16	1.12	$>$	1.10	0.97	$<$	1.06	1.03	$<$	1.09	0.95	$<$	1.21

3.4. Results: "Perfect" Labor Market Conditions and "Reduced Mobility"

This section presents additional simulations to answer two additional questions: First, what happens under perfect labor market conditions, i.e. when each applicant applies to each open position. Second, does the mechanism of statistical discrimination intensify the penalty of a group when this group is less mobile on the market, or has limited information about open positions. Both reduced mobility and limited information on attractive labor market positions translate into fewer applications of this group to employers in the simulation model (the opposite of "perfect" labor markets).

In the following runs, therefore, the number of applications for both groups varies between the above used 10 applications and the number of open positions, which was set to 160 in all runs. More exactly the simulation was set to generate 10, 60, 110, or 160 random employer numbers (see section 3.1 for details). However, as described in section 3.1.2, it may happen that a worker "randomly selects" the same employer number more than once. Therefore, the actual unique applications are below the intended numbers. Empirically on average 9.7 (instead of 10), 50.2 (instead of 60), 79.7 (instead of 110) and 101.3 (instead of 160) unique applications have been simulated. The deviation from the intended number of applications increases, as the likelihood of getting identical employer numbers raises the more often one draws from the same universe of employers. The simulations is applied on a share of 20% Bs of in total 500 workers, a stable γ_W of 0.8 and again γ_B varies between 0.2 and 0.8.

[figure 4 about here]

Figure 4 presents average odds ratios (see above) for the simulated parameter space. The basic outcome is the same, regardless of the additional variation in average number of applications per worker. The group whose signals are trusted less, is discriminated more the higher the perceived reliability differential is. In figure 4 all odds ratios rise way above 1 the larger the reliability differential between Ws and Bs gets (1 would indicate equality, values > 1 indicate an advantage for Ws). The solid black line indicating a perfect labor market, in which both groups apply to all open positions, is not systematically different from the bluish-gray dash-dot line that indicates all workers to apply to ten positions only (in all other simulations presented so far, 10 applications per worker were used). Essentially all lines indicating equal numbers of applications by Ws and Bs result in the same pattern and are even somewhat difficult to distinguish. More interesting are the lines from runs where the number of applications is different between Ws and Bs. The highest levels of inequality result when the minority group B applies only for 10 positions and the Ws apply for more positions (see the bluish-gray solid, dash and dot lines). On the opposite, the lowest levels are predicted when Ws only apply for 10 positions and Bs apply for more (see the lowest three dash-dot lines). In sum, we observe the statistical discrimination mechanism to be robust against different levels of applications. Even under the unrealistic assumption that there is a "perfect" labor market, i.e. every worker applies for every open position, statistical discrimination remains. Second, the effect is amplified when one group applies to fewer open positions than the other. The employment rates of a minority group, whose signals are less trusted, will be even lower when this group is less mobile, and hence applies to fewer employers.

3.5. Robustness Check: More Variance in Parameters

This section will try to answer the question whether the basic findings hold, when more variance at one time is allowed on all parameters (e.g. variation in γ, varying number of applications, ...)? I also wonder if it makes a difference when instead of continuous a discrete signal is used. E.g.

school/university leaving certificates or vocational degrees rather generates a very limited set of discrete gradations...

[noch nicht fertig]

4. Conclusion

The paper focussed on whether or not Phelp's measurement model of statistical discrimination can explain inequality in hiring, i.e. group discrimination against workers whose productivity signals are perceived on average as less reliable. I used the example of ethnic inequality, where it is usually assumed that an ethnic minority's signals are trusted less, whereas employers perceive the majority's signals as more reliable. The *theoretical analysis* finds that in a singular hiring situation one "majority" applicant with a signal above the mean of productivity signals would be sufficient to create an advantage for the majority group. If we consider a whole labor market the analysis shows that the direction of discrimination depends on the relation of workers seeking a job to the number of open positions. A minority group whose signals are trusted less is discriminated in very competitive labor markets, whereas, under the conditions of less competition the model even predicts discrimination against those workers whose signals are trusted more. Only under the very specific condition of the relation of open positions to applicants being 1:2 the Phelp's model does not result in inequality. Considering the access to *qualified positions* the model always results in discrimination against the group whose signals are trusted less, regardless of the level of competition. This is true for all qualified positions that require a skill level above the average productivity in the market (a very realistic assumption). *Simulations*, in lack of sufficient empirical data, show that all theoretically derived effects occur, and that they are solely based on group differences in perceived reliability. Further analyses show that reduced mobility, translated into fewer applications to employers in the simulations, will pronounce these effects substantially. *To sum up*, Phelp's measurement model of statistical discrimination results in group discrimination in almost all structural conditions. However, very against intuition, the predicted discrimination effects do not always reduce the minority group's labor market attainment (whose signals are perceived less trustworthy). Only when there is high competition about positions, or when we consider inequality in job quality, Phelp's model can account for ethnic disadvantages.

To conclude inequality research should take statistical discrimination into account more seriously. Even if the original Phelp's models comes under certain conditions to somewhat unrealistic predictions, it has a couple of advantages to offer. The unrealistic predictions to some extent can be attributed to very idealistic assumptions concerning employers: The original model, while convincingly taking into account the bounded rationality of employers in regard to the reliability of signals, is still very demanding in respect to the knowledge employers are assumed to have concerning average productivity. It requires employers to possess up-to-date undisturbed beliefs about the average productivity of groups of workers. Aside from this "flaw" the basic mechanism is precise yet simple and to some extent it takes into account the bounded rationality of employers very realistically. Furthermore, it offers "interfaces" to include competing discrimination models. For example, Becker's tastes (Becker 1971: 40) or status based discrimination theory that focusses on cognitive stereotypes about productivity, based on ascriptive categories (e.g. Correll und Benard 2006) or e.g. "error discrimination" (England und Lewin 1989: 242) could be modeled into group-specific α-parameters. Differences in group-specific γ values originally result from the minority status and the following fewer experiences of employers with minority group applicants. However, in principal it would also allow to model "social resources" as network effects or personal recommendations ([1]; Lin 1999). E.g. personal recommendations should raise the trust into individual signals. *In sum*, to come to an inclusive and realistic theoretical model of employer behavior that can explain discriminative hiring, the measurement model of statistical discrimination offers the most promising starting-point.

References

[1] Mark Granovetter. The strength of weak ties. *American Journal of Sociology*, 78:1360–1380, 1973.

Figure 2: Simulation of Access to Positions

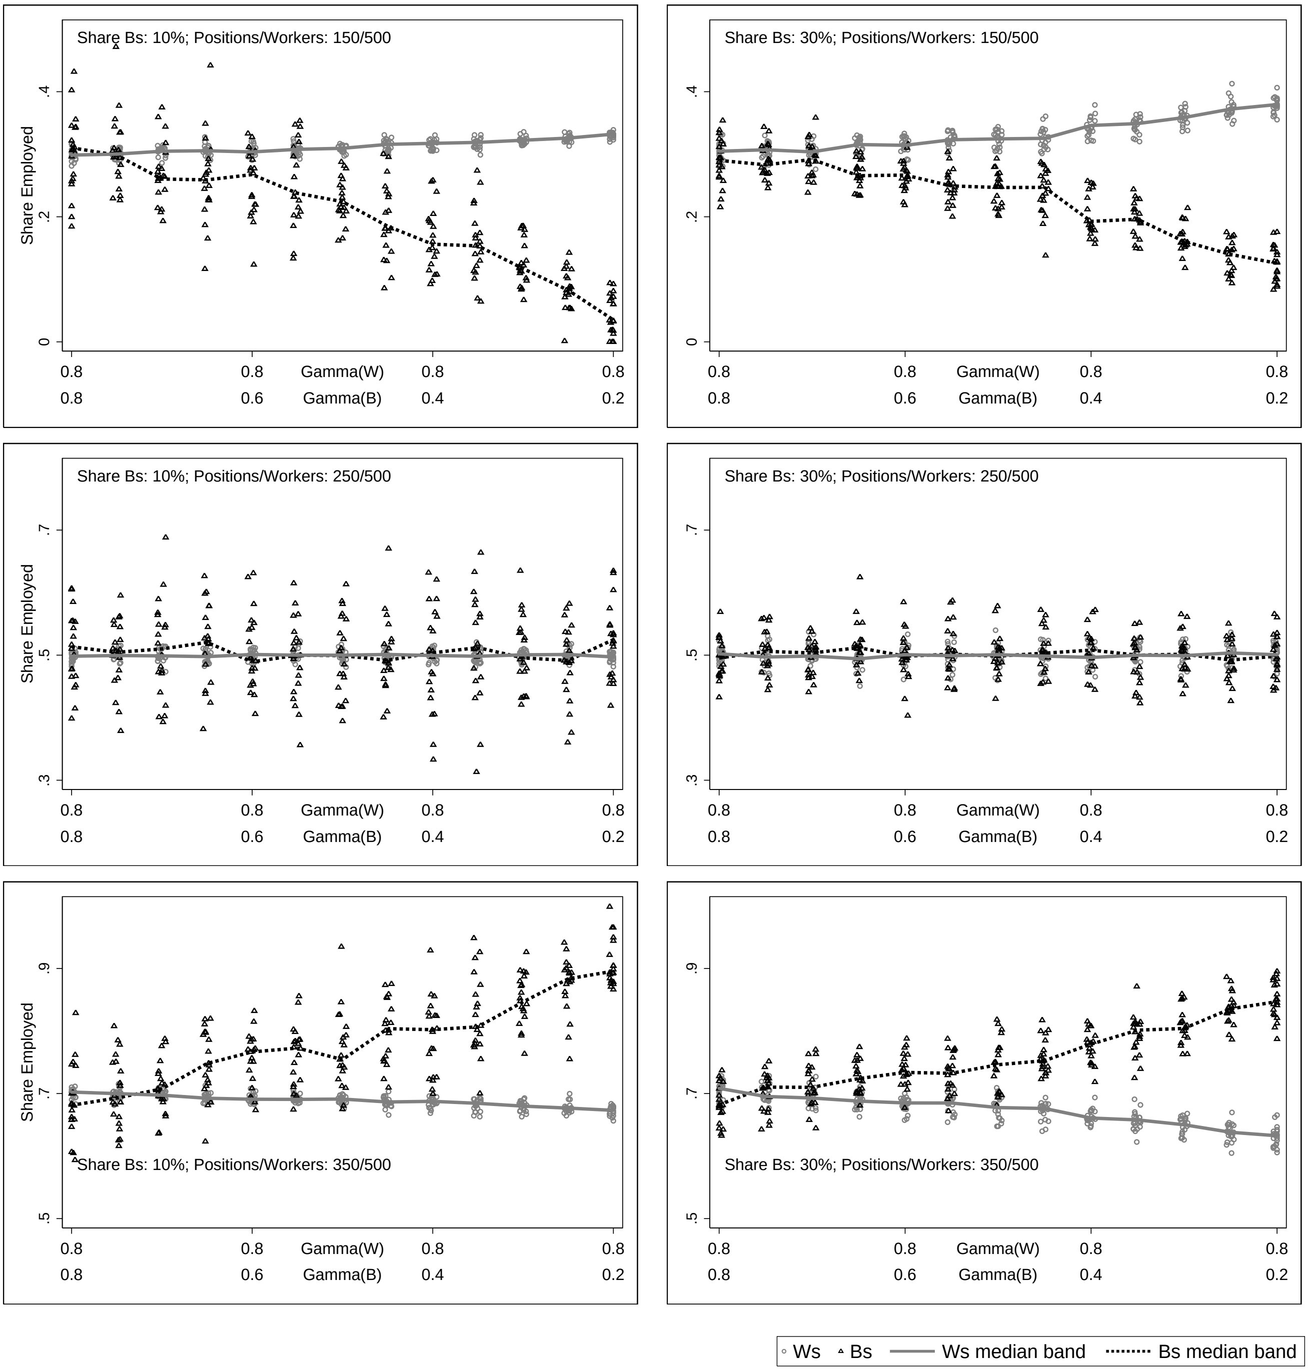

Figure 3: Simulation of Access to Qualified Positions

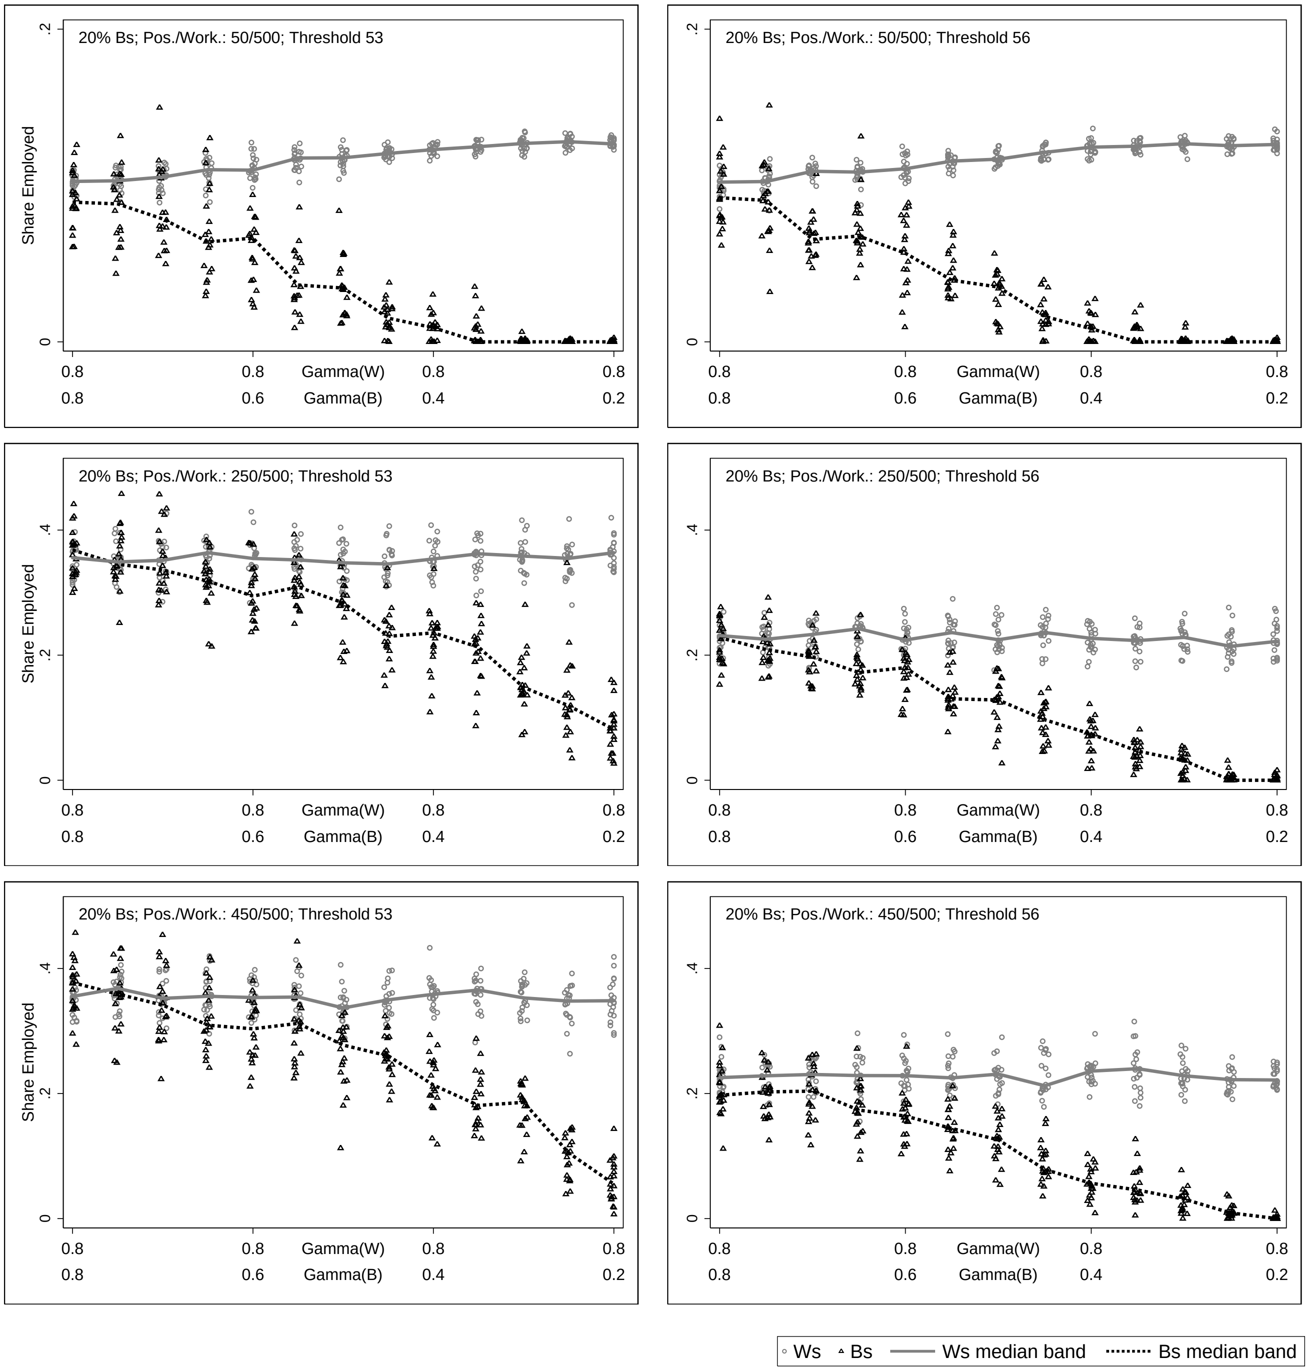

Figure 4: "Perfect" Labor Market Conditions and "Reduced Mobility"

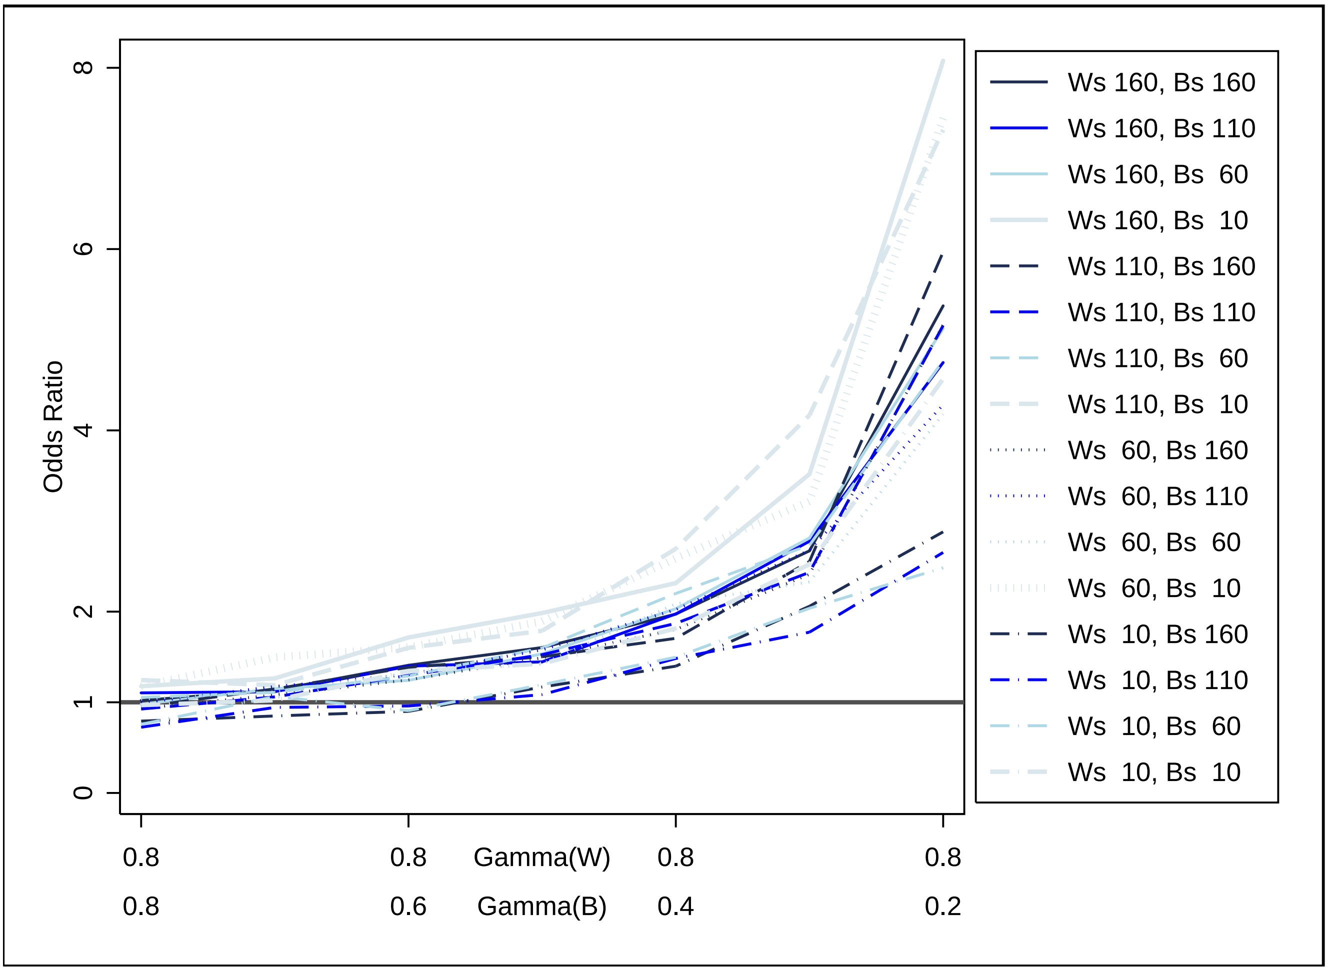

A. A Basic Check of the Simulation Model (Appendix)

OLS-Regression of Share of B-Workers employed

VARIABLES	(1) Model_1	(2) Model_2
γ_W - γ_B	-0.233***	-0.230***
	(0.007)	(0.004)
W Workers: Mean Signal		-0.014***
		(0.001)
B Workers: Mean Signal		0.016***
		(0.001)
W Workers: Std. Dev. Signal		-0.011***
		(0.002)
B Workers: Std. Dev. Signal		0.009***
		(0.002)
Constant	0.209***	0.115
	(0.002)	(0.094)
Observations	260	260
R^2	0.826	0.932

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

OLS-Regression of Share of W-Workers employed

VARIABLES	(1) Model_1	(2) Model_2
γ_W - γ_B	0.238***	0.235***
	(0.007)	(0.004)
W Workers: Mean Signal		0.014***
		(0.001)
B Workers: Mean Signal		-0.016***
		(0.001)
W Workers: Std. Dev. Signal		0.012***
		(0.002)
B Workers: Std. Dev. Signal		-0.008***
		(0.002)
Constant	0.190***	0.243**
	(0.002)	(0.096)
Observations	260	260
R^2	0.833	0.931

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

B. Simulation Code (Stata)

```
******************************************************************************/*

      S i m u l a t e   S t a t i s t i c a l   D i s c r i m i n a t i o n

                            Christian Hunkler
                    Mannheim University, August 2009

******************************************************************************/
*>> CREATE WORKERS (workers in rows; first 5rows reserved for employer param.)
    clear
    global obs = $l + 6
    set obs $obs
    gen id = _n // for sorting of datasets

    *>> Generate workers Signals (s) and allot to two groups (b)
    if $sames==1 {
        qui gen s =  rnormal($sm,$sd)  if id>6
    }

    gen     b = int(runiform()*100)  if id>6
    replace b = 1 if b <=$b          &  id>6
    replace b = 0 if b > $b          &  id>6

*>> CREATE EMPLOYERS (in columns)
    local n = 1
    while ‘n’ <= $e {
        quietly generate e‘n’ = . if id>6
        quietly local n = ‘n’ + 1
    }

    *>> write employers’ gammas in lines 1 and 2
    gen str50 desc=""
    replace   desc="Gamma W" if id==1
    replace   desc="Gamma B" if id==2
    replace   desc="Alpha W" if id==3
    replace   desc="Alpha B" if id==4
    order desc

    gen tempgw = rnormal($gw, $gwsd)  // temporary help variable for gamma
    gen tempgb = rnormal($gb, $gbsd)

    sum s if id>6
    local alphab=r(mean)
    local alphaw=r(mean)

    local n = 1
    while ‘n’ <= $e { // go through all employers
        local gw = tempgw in ‘n’
        local gb = tempgb in ‘n’
        quietly replace e‘n’= ‘gw’/100 if id==1 // this is gamma w
        quietly replace e‘n’= ‘gb’/100 if id==2 // this is gamma b
        quietly replace e‘n’= ‘alphaw’ if id==3 // this is alpha w
        quietly replace e‘n’= ‘alphab’ if id==4 // this is alpha b
        quietly local n = ‘n’ + 1
    }
    capture drop temp*

*>> APPLICATION
    local n = 1     //write variables with randomly selected employer number
    while ‘n’ <= $aw {
        gen apply_w‘n’ = int(runiform()*$e) + 1  if id>6
        local n = ‘n’ + 1
    }

    local n = 1
    while ‘n’ <= $ab {
        gen apply_b‘n’ = int(runiform()*$e) + 1  if id>6
        local n = ‘n’ + 1
    }

    forvalues i=1(1)$e { // copy signals of applications into employer columns
        forvalues j=1(1)$aw {
            quietly replace e‘i’ = s if b==0 & apply_w‘j’==‘i’
        }
    }

    forvalues i=1(1)$e {
        forvalues j=1(1)$ab {
            quietly replace e‘i’ = s if b==1 & apply_b‘j’==‘i’
        }
    }

    *>> Output Variable: number of applications by Ws and Bs for each employer
    forvalues i=1(1)$e {                        // single constant var for each e
        quietly gen     temp‘i’= 1 if e‘i’!=. & b==0
        quietly egen    appl_w_‘i’ = total(temp‘i’)
        quietly drop temp‘i’
        }

    gen appl_w = .           // write into one variable
    forvalues i=1(1)$e {
        quietly replace appl_w = appl_w_‘i’ in ‘i’
        quietly drop appl_w_‘i’
    }
```

```
    forvalues i=1(1)$e {                        // single constant var for each e
        quietly gen      temp‘i’= 1 if e‘i’!=. & b==1
        quietly egen     appl_b_‘i’ = total(temp‘i’)
        quietly drop temp‘i’
        }

    gen appl_b = .            // write into one variable
    forvalues i=1(1)$e {
        quietly replace appl_b = appl_b_‘i’ in ‘i’
        quietly drop appl_b_‘i’
    }


    *>> HIRING ACCORDING TO STATISTICAL DISCRIMINATION
    capture drop employed            // help variable
    gen employed = 0 if id>6

    forvalues i=1(1)$e {
        quietly sum e‘i’ in 1        // store employers parameters as locals
        local gw = r(mean)
        quietly sum e‘i’ in 2
        local gb = r(mean)
        quietly sum e‘i’ in 3
        local aw = r(mean)
        quietly sum e‘i’ in 4
        local ab = r(mean)

        quietly gen esd‘i’= e‘i’    // generate copy of signal of applicants

        quietly replace esd‘i’ = ((1-‘gw’)*‘aw’) + ( ‘gw’ * esd‘i’) if b==0
        quietly replace esd‘i’ = ((1-‘gb’)*‘ab’) + ( ‘gb’ * esd‘i’) if b==1

        quietly sum esd‘i’            if employed==0 & id>6 // search best appl.

      if $qualified==0 { // Simulation of access to positions. no thresholds
        quietly replace esd‘i’   = 0 if esd‘i’!=r(max) & esd‘i’!=. & id>6
        quietly replace employed = 1 if esd‘i’==r(max) & esd‘i’!=. & id>6
      }
      if $qualified==1 { //Simulation of access to qualified pos.: thresholds
        qui replace esd‘i’  =0 if esd‘i’!=r(max) &esd‘i’!=. &id>6 |esd‘i’< $th
        qui replace employed=1 if esd‘i’==r(max) &esd‘i’!=. &id>6 &esd‘i’>=$th
      }

        move esd‘i’ e‘i’

        if $check1==1  { // Extra Output to trace queuing of workers
            preserve
                keep if id>6
                keep b s employed /* esd‘i’ */
                saveold "$ssdout/$name/${name}_check1_${run}_‘i’.dta"
            restore
        }
    }

    if $check1==1 {       // Extra Output to trace queuing of workers
        preserve
            noisily use     "$ssdout/${name}_check1_${run}_1.dta", clear
            gen e=1
            erase   "$ssdout/$name/${name}_check1_${run}_1.dta"
            forvalues i=2(1)$e {
                append using "$ssdout/$name/${name}_check1_${run}_‘i’.dta"
                replace e=‘i’ if e==.
                erase        "$ssdout/$name/${name}_check1_${run}_‘i’.dta"
            }
            compress
            saveold "$ssdout/$name/${name}_check1_${run}.dta", replace
        restore

    }

    capture drop employed

    *>> HIRING ACCORDING TO HIGHEST SIGNAL
    capture drop employed
    gen employed = 0 if id>6

    forvalues i=1(1)$e {
        quietly sum e‘i’ if employed==0 & id>6
      if $qualified==0 { // Simulation of access to positions. no thresholds
        quietly replace e‘i’     = 0 if e‘i’!=r(max) & e‘i’!=. & id>6
        quietly replace employed = 1 if e‘i’==r(max) & e‘i’!=. & id>6
      }
      if $qualified==1 { //Simulation of access to qualified pos.: thresholds
        qui replace e‘i’     = 0 if e‘i’!=r(max) &e‘i’!=. &id>6 |e‘i’< $th
        qui replace employed = 1 if e‘i’==r(max) &e‘i’!=. &id>6 &e‘i’>=$th
      }
    }
    capture drop employed

*>> CONSTRUCT OUTCOME VARIABLES (EMPLOYED):
    gen empls  = 0
    gen emplsd = 0

    forvalues i=1(1)$e {
        quietly replace empls  = empls  + 1 if e‘i’  >0 & e‘i’  <. & id>6
        quietly replace emplsd = emplsd + 1 if esd‘i’>0 & esd‘i’<. & id>6
    }


*>> WRITE FINAL OUTPUT DATASETS
```

```
    compress
    saveold "$ssdout/$name/${name}_$run.dta"        // SAVE MIKRO-FILES

    gen start = .                                    // SAVE MAKRO-FILES (Start)
    order start
    gen stop = .

    gen run      = $run
    gen seed     = c(seed)
    gen l        = $l
    gen e        = $e
    gen bs       = $b

    gen gw       = $gw
    gen gwsd     = $gwsd
    gen gb       = $gb
    gen gbsd     = $gbsd
    gen aw       = $aw
    gen ab       = $ab
    if $qualified==1 { // Simulation of access to positions: with thresholds
        gen th = $th
    }

    foreach var of varlist s b empls emplsd appl_w appl_b {
        quietly sum ‘var’          if id>6
        gen ‘var’_sd     = r(sd)
        gen ‘var’_mean   = r(mean)
        gen ‘var’_min    = r(min)
        gen ‘var’_max    = r(max)
    }
    foreach var of varlist s b empls emplsd {
        quietly sum ‘var’ if b==0 & id>6
        gen ‘var’_w_sd     = r(sd)
        gen ‘var’_w_mean   = r(mean)
        gen ‘var’_w_min    = r(min)
        gen ‘var’_w_max    = r(max)
    }
    foreach var of varlist s b empls emplsd {
        quietly sum ‘var’ if b==1 & id>6
        gen ‘var’_b_sd     = r(sd)
        gen ‘var’_b_mean   = r(mean)
        gen ‘var’_b_min    = r(min)
        gen ‘var’_b_max    = r(max)
    }

    drop start-stop // mikro variables
    keep in 1       // only aggregate findings exported for makro dataset)
    compress
    saveold "$ssdout/$name/${name}_${run}_makro.dta"

 exit

*******************************************************************************
```