

DOCUMENT DE TRAVAIL

N° 12-2012

First Order Stochastic Dominance and the Measurement of Hiring
Discrimination: A ranking extension of correspondence
testings with an application to gender and origin

Emmanuel Duguet, Loïc Du Parquet, Yannick L'Horty, Pascale Petit

First Order Stochastic Dominance and the Measurement of Hiring Discrimination: A ranking extension of correspondence testings with an application to gender and origin

Emmanuel Duguet

Université Paris Est, ERUDITE and TEPP (FR CNRS 3435)

Loïc Du Parquet

Université du Maine, GAINS and TEPP (FR CNRS 3435)

Yannick L'Horty

Université Paris Est, ERUDITE and TEPP (FR CNRS 3435)

Pascale Petit

Université d'Evry Val d'Essonne and TEPP (FR CNRS 3435)

August, 2012

Abstract - We extend the standard hiring discrimination measure by including the cases where several candidates are invited to the same interview. The new measure considers the order in which the employer will contact the candidates as opposed to considering only whether or not a job applicant is invited to an interview – a practice common in the previous literature. We propose to apply the first order stochastic dominance (FOSD) criterion to the ranking of the candidates, which appears to be especially relevant for hiring discrimination. We show theoretically that FOSD always implies a positive value for the standard discrimination coefficient used in the literature, and that the converse is false. We apply our analysis to a correspondence testing that has been conducted in the Paris region. We sent 8 fictitious candidates with a Master's degree to the same 310 job offers in computing in order to measure gender and origin discrimination. We found that - out of 28 possible comparisons - there are 25 cases of stochastic dominance that we interpret as strong discrimination against some candidates. In our application, the standard discrimination coefficient tends to underestimate the degree of discrimination.

JEL classification: C93, J16, J61, J71.

Keywords: gender, origin, hiring discrimination, first order stochastic dominance.

Corresponding author: Emmanuel Duguet, Université Paris-Est, ERUDITE, Mail des Mèches, 61 avenue du Général de Gaulle, 94010 Créteil cedex, France, emmanuel.duguet@u-pec.fr

1. Introduction

Correspondence testing is an appropriate technique for evaluating hiring discrimination in the labor market. Economists have been using this methodology since the end of the seventies (Firth, 1981, 1982). The technique is a controlled experiment that consists of fabricating written applications (CVs and cover letters) of fictitious candidates. The applications are identical, except for one characteristic *a priori* not tied to productivity, like the applicant's gender or origin. The applications are sent to the same job offers, and one examines whether the applicants obtained a similar access to hiring interviews.

RIACH and RICH (2002) highlighted two trends emerging from the results of the correspondence testings conducted in order to evaluate the degree of gender discrimination in hiring. First, women suffer from discrimination in hiring for well-paid positions involving responsibilities. In Great Britain, for instance, FIRTH (1982) found that women were less likely than men to obtain a position as a qualified accountant; in Philadelphia, NEUMARK et al. (1996) found significant discrimination against women for server jobs in high-price restaurants and significant discrimination against men for server jobs in low-price restaurants. Second, discrimination can be observed in jobs in which one gender is over-represented. For example, there is discrimination against men in secretary jobs traditionally occupied by women; conversely, there is discrimination in hiring against women in activities traditionally occupied by men, like mechanics (WEICHSELBAUMER, 2004). Lastly, BERTRAND and MULLAINATHAN (2003) stressed that the scale of discrimination in hiring varies according to the ethnic origin of the candidates. Their testing evaluated the scale of racial discrimination according to the gender of the job applicants. The authors found that candidates with Caucasian-sounding names receive 50% more callbacks than their counterpart with African-American sounding names.

From a methodological viewpoint, the correspondence testing literature consists in sending resumes to a recruiter that invites the candidates to an interview or not. Consider the case of two candidates, A and B. Four answering cases are possible: no candidate is invited, only candidate A is invited, only candidate B is invited or both candidates are invited. The standard discrimination measure is a calculation of the following discrimination coefficient. One computes the difference between the percentage of offers for which candidate A had been invited without candidate B and the percentage of offers for which candidate B had been invited without candidate A. With this method, the cases where both candidates had been invited are not used in the measurement of discrimination because both candidates are considered to have been equally treated.

Our aim is to provide a more general measure of hiring discrimination. We claim that the cases in which both candidates have been invited should also be included in the analysis when information is available about the order in which the candidates have been called. The justification is that in the standard testing procedures, the fictitious candidates have instructions to respond that they have already found a job when they are called, so that they always decline the job offer. Therefore, the recruiters decide to go further down the short list of candidates, and thus reveal a ranking of the candidates. In this paper we propose using the condition of first order stochastic dominance (henceforth, FOSD) on the candidates' ranking. Consider the case with $k \geq 3$ candidates. Candidate A FOSD candidate B when he/she has higher probabilities to rank first than B, to rank among the first two candidates than B, to rank among the first three candidates than B etc. We interpret the condition of FOSD as the existence of strong discrimination against candidate B in this paper, because we show that FOSD involves a stronger form of discrimination than the one measured by the standard discrimination coefficient.

We apply this new method to a correspondence testing conducted in 2009 which aimed to measure both gender and ethnic origin discrimination among French young workers in the Paris region. We fabricated eight similar resumé: four origins for each gender (French, Moroccan, Senegalese, and Vietnamese). The origin is indicated by the sound of the first and last names. The eight candidates all hold French nationality, live in Paris, and apply for the same jobs as software developers at the Master's degree level. Applications have been sent simultaneously to the same 310 job offers. This protocol enables us to examine first, whether hiring discrimination is based on origin, and second, whether discrimination against women varies with their origin.

We show that, while the standard discrimination coefficients point in the right direction in most of the cases, they tend to underestimate the importance of discrimination in the labor market. The next section presents the methodology, and the following presents the application and summarizes our results.

2. Methodology

2.1 Standard discrimination coefficient

Consider a recruiter with preferences for the candidates A and B represented by the utilities u_A and u_B . These utilities are specific to each recruiter and result from pre-conceptions about the candidates, because the candidates are equally productive by construction of the testing experiment.

Each recruiter has a reservation utility level v_R above which the candidates are invited to an interview. We define the relative utility levels $u_A = v_A - v_R$ and $u_B = v_B - v_R$. The four potential answering cases can be represented in the following way. If $u_A < 0$ and $u_B < 0$, no candidate is invited to an interview. When $u_A < 0 < u_B$, only candidate B is invited; when $u_B < 0 < u_A$, only candidate A is invited. Finally, when $u_A > 0$ and $u_B > 0$, both candidates are. These cases are illustrated in Figure 1. The standard measure of discrimination against candidate B, used in the literature, considers only cases in which only one of the two candidates is invited. These cases are illustrated by the North-West and the South-East quadrants of Figure 1. We denote this discrimination coefficient as $\Delta_1(A, B)$:

$$\Delta_1(A, B) = P[u_B < 0 < u_A] - P[u_A < 0 < u_B] = P[A \text{ invited, } B \text{ uninvited}] - P[B \text{ invited, } A \text{ uninvited}]$$


According to this measure, there is no discrimination when both candidates have equal chances to be invited, and a positive number indicates that candidate A is – on average – preferred to candidate B.¹ Also notice the property that $\Delta_1(A, B) = -\Delta_1(B, A)$.

2.2 First order stochastic dominance and strong discrimination

We propose to extend the standard measure of discrimination Δ_1 to the ranking of the candidates when both are invited, which is equivalent to consider *all* the quadrants in Figure 1.

Figure 1: Interview decisions (gender case)

¹ Since Δ_1 measures discrimination when only one of the two candidates is invited to the interview, we use the subscript 1.



An important point to mention is that we need to consider the ranking among $k \geq 3$ candidates, since our application includes 8 candidates. Therefore we develop our analysis in the general case of an undetermined number of candidates k . Among these k candidates, we perform all the pairwise

comparisons available (here : $\frac{k(k-1)}{2} = 28$).

In order to compare the rankings of two candidates, we use the concept of first order stochastic dominance. Suppose that the k candidates are ranked according to the recruiter's utilities. The candidates that have not been invited satisfy the condition that $u_j = v_j - v_R < 0$. The ranking of the candidates (from 1st to k^{th}) results from the recruiter's tastes for the candidates. The highest utility corresponds to the candidate ranked first, and negative utilities correspond to the candidates that have not been invited. In order to perform our analysis, we need to separate the candidates that have not been invited from the others by creating a rank $k+1$. This additional rank is required because the candidates that have not been called cannot be ranked between themselves. This also allows us to defined A_1 directly. We know only that the uninvited candidates' utilities are below the recruiter's reservation utility levels and therefore that they are ranked behind the candidates that have been invited.

Consider first the case for which all the candidates have been invited. Using the order statistic², we obtain the ranking of the utilities of the candidates, $0 < u_{(1)} \leq u_{(2)} \leq \dots \leq u_{(k)}$, that corresponds to the ranking $k, k-1, \dots, 1$. When only j candidates are invited, we have the ranking $u_{(1)} \leq \dots \leq u_{(k-j)} < 0 < u_{(k-j+1)} \leq \dots \leq u_{(k)}$ that corresponds to the ranking $k+1, \dots, k+1, j, \dots, 1$.

The first stochastic dominance of candidate A over candidate B is defined as:

$$\Pr[u_A \geq u] \geq \Pr[u_B \geq u] \quad \forall u, \text{ and } \exists \bar{u} \text{ such that } \Pr[u_A \geq \bar{u}] > \Pr[u_B \geq \bar{u}],$$

which means that candidate A has a higher probability to reach a given utility level than candidate B. This relationship is especially easy to interpret when u is set at the reservation utility level of the recruiter, since it means that candidate A has a higher probability to be invited to the interview than candidate B. We also see that FOSD covers more cases than the standard discrimination measure because it makes use of all possible utility thresholds.

For practical reasons we will work with the ranks, since they are observable while the utilities are not. We just need to reverse the inequalities inside the probabilities, since the higher the utility the lower the rank (rank 1 for the most preferred candidate with utility $u_1((k))$):

$$\Pr[r_A \leq r] \geq \Pr[r_B \leq r] \quad \forall r \in \{1, \dots, k+1\} \text{ and } \exists \bar{r} \text{ such that } \Pr[r_A \leq \bar{r}] > \Pr[r_B \leq \bar{r}]$$

Consider the case $r = 1$. Then $\Pr[r_A \leq 1] = \Pr[r_A = 1]$, which gives the probability to be ranked first. If the corresponding inequality holds, the candidate A has a higher probability to be ranked first than the candidate B. Now set $r = 2$. We conclude that candidate A has a higher probability to be ranked among the two first candidates than candidate B. Performing the comparisons up to $r = k$, $\Pr[r_A \leq k]$ is the probability that candidate A is invited to the interview. Therefore candidate A has a higher probability to be invited to the interview than candidate B. In summary, when A FOSD B, the candidate A always has a higher probability to be in the leading group than the candidate B, whatever the definition of the leading group is. This definition is especially relevant for the measurement of discrimination, and this is what motivates our use of FOSD. Graphically, FOSD means that the CDF of candidate A – defined on ranks – stands above the CDF of candidate B.³

² The order statistic sorts the utility values in ascending order so that the utility vector (u_1, \dots, u_k) becomes $(u_{(1)}, \dots, u_{(k)})$ where $u_{(1)}$ is the smallest utility and $u_{(k)}$ is the highest utility.

³ The CDF defined over the ranks is equivalent to the survival function defined on the utilities.

We demonstrate the following property for the comparison of two candidates, A and B, in a testing with k candidates.

Property:

$$(a) \ A \text{ FOSD } B \Rightarrow \Pr[A \text{ invited}, B \text{ uninvited}] \geq \Pr[B \text{ invited}, A \text{ uninvited}],$$

or equivalently : $A \text{ FOSD } B \Rightarrow \Delta_1(A, B) \geq 0$.

(b) The converse of this implication is false.

In order to simplify the exposition, we introduce the following notation, considering candidates A and B among k candidates:

$$\Pr(r_A = a \cap r_B = b) = \Pr(a, b),$$

and we use the two following conventions:

- 1) Each of the k candidates has $k+1$ possible ranks: $1^{\text{st}}, 2^{\text{nd}}, \dots, k^{\text{th}}$ or *uninvited* which we denote $(k+1)^{\text{th}}$.
- 2) When there are no ties, we set $\Pr(a, a) = 0$, which means that two candidates cannot have the same rank. The existence of ties does not affect the property.

On the other hand, it is always true that $\Pr(r_i \leq r) = \sum_{m=1}^r \Pr(r_i = m)$, $i \in \{A, B\}$,

and $\Pr(r_i = m) = \sum_{j=1}^{k+1} \Pr(m, j)$, so that $\Pr(r_i \leq r) = \sum_{m=1}^r \sum_{j=1}^{k+1} \Pr(m, j)$.

Proof:

(a) Notice that $\Pr[r_i \leq k]$ is the probability that candidate i is invited, regardless of the situation of the other candidates. FOSD implies that:

$$\Pr[r_A \leq k] \geq \Pr[r_B \leq k]$$

$$\Leftrightarrow \sum_{m=1}^k \sum_{j=1}^{k+1} \Pr(m, j) \geq \sum_{m=1}^k \sum_{j=1}^{k+1} \Pr(j, m)$$

$$\Leftrightarrow \sum_{m=1}^k \sum_{j=1}^k \Pr(m, j) + \sum_{m=1}^k \Pr(m, k+1) \geq \sum_{m=1}^k \sum_{j=1}^k \Pr(j, m) + \sum_{m=1}^k \Pr(k+1, m) \quad (1),$$

and the sum $\sum_{m=1}^k \sum_{j=1}^k \Pr(m, j)$ adds the probabilities of all the possible couples (m, j) between 1 and k . Therefore :

$$\sum_{m=1}^k \sum_{j=1}^k \Pr(m, j) = \sum_{m=1}^k \sum_{j=1}^k \Pr(j, m) \quad \text{and inequality (1) simplifies to:}$$

$$\sum_{m=1}^k \Pr(m, k+1) \geq \sum_{m=1}^k \Pr(k+1, m)$$

$$\Leftrightarrow \Pr[\text{A invited, B uninvited}] \geq \Pr[\text{B invited, A uninvited}]$$

$$\Leftrightarrow \Delta_1(\text{A}, \text{B}) \geq 0$$

(b) The second part of the proof is easier. We have shown in part (a) that $\Delta_1(\text{A}, \text{B}) \geq 0$ is a consequence of the last inequality that defines FOSD, and therefore it does not imply FOSD.

Q.E.D.

This allows us to define strong discrimination against candidate B as A FOSD B, and weak discrimination as $\Delta_1(\text{A}, \text{B}) \geq 0$ or equivalently $\Delta_1(\text{B}, \text{A}) \leq 0$. Weak discrimination against candidate B simply states that the probability to be invited to an interview alone is higher for candidate A than for candidate B, while strong discrimination implies that the probability that candidate A has a good ranking is higher than for candidate B. Therefore weak discrimination uses only the extreme case where one candidate only is invited, while strong discrimination also includes all the cases where both candidates are invited. This has three interesting consequences.

First, two candidates with similar values for Δ_1 might or might not display FOSD. This opens the possibility to classify candidates with similar discrimination coefficients according to the FOSD criterion. Second, the importance of Δ_1 is not always a measure of the importance of discrimination, because a candidate experiencing strong discrimination can have a lower discrimination coefficient than one experiencing weak discrimination. We find evidence for both properties in our application. Third, it is possible to provide an overall ranking of the discrimination faced by the candidates by using the conditions of weak and strong discrimination together. We will first rank the candidates according to the strong discrimination criterion and, when they do not FOSD each other, we will rank

them according to the weak discrimination criterion. Combining these two measures, we are able to indicate each candidate's position on an ordinal discrimination scale.

Figures 2 and 3 provide two examples of the presence or absence of FOSD taken from our data. In Figure 2, we compare the Vietnamese origin candidates, and find that women dominate men. But it is also possible that two candidates do not dominate each other; this happens when one CDF crosses the other. In Figure 3, we compare the French origin candidates and find that none dominates the other. The CDFs of all the candidates are presented in Table 4, and the FOSD analysis in Tables 5A to 5C (28 pairwise comparisons).

Figure 2: Vietnamese origin women FOSD Vietnamese origin men





Figure 3: French origin candidates do not FOSD each other



3. Application and results

We performed a correspondence testing in the Paris region. It has been designed to assess the effects of gender and ethnic origin on access to hiring interviews. Eight CVs of fictitious candidates with similar characteristics regarding productivity were constructed for young software developers with Masters' degrees. These candidates, all of French nationality, differed solely by gender and ethnic origin indicated by their names (French, Moroccan, Senegalese or Vietnamese).

3.1 Data collection

This stage was designed with a few challenges in mind. First, we wanted to limit the probability of detection when sending eight CVs simultaneously. Second, we also needed to consider a job with a high labor demand in order to have high response rates. This methodological precaution proved to be particularly useful in a context of economic recession. For this purpose, we consulted the Historic Data File of the French unemployment agency (Pôle Emploi) in order to find a profession where both the number of candidates and the labor demand was high. The profession that met these criteria is software development.

The 8 fictitious candidates explicitly state their French nationality on their CVs; their first and last names indicate their gender and ethnic origin. The first names given to the candidates are common, given the origins of the candidates, and the last names associated with each origin are among the most widespread. The candidates live in inner Paris, and in similar districts in terms of geographic position and demographic composition. The individual characteristics of the 8 fictitious candidates are presented in Table 1. In the case of the Vietnamese origin candidates, we had to indicate the gender explicitly on the CV, because the first names (Tien Hiep and Minh Trang) were not widespread enough to identify gender.

Table 1: Characteristics of the 8 candidates

Candidates for job applications	Gender	Ethnic origin
Anthony BERTRAND <i>Place of residence: Paris 14th district</i>	Male	French
Sophie MOREAU <i>Place of residence: Paris 14th district</i>	Female	
Abdallah ZALEGH <i>Place of residence: Paris 13th district</i>	Male	Moroccan
Jamila KAIDI <i>Place of residence: Paris 14th district</i>	Female	
Amadou DIALLO <i>Place of residence: Paris 13th district</i>	Male	Senegalese
Fatou DIOUF <i>Place of residence: Paris 13th district</i>	Female	
Tien Hiep PHAM <i>Place of residence: Paris 13th district</i>	Male (indicated on CV)	Vietnamese
Minh Trang TRAN <i>Place of residence: Paris 13th district</i>	Female (indicated on CV)	

The applications were sent in response to the same job offers, and were equivalent in terms of productivity-related characteristics. The candidates possessed identical diplomas, professional careers and experience, and the same computing and language skills. None of the candidates had been unemployed, and they all had jobs at the time of their applications. In addition, we had these fictitious CVs checked by recognized professional experts in the field. More details on the experimental design are given in Appendix 1.

Because these applications were sent at the same time in response to the same job offers, they had to include some elements of differentiation. These differences involved the presentation of the CV: font type and size, page layout, while still remaining standard. The candidates declared professional experience acquired in real companies, which were different but comparable (in terms of activity, size and market power). The candidates hobbies were also different, while remaining very standard and impersonal (sport, cinema, reading, music, etc.). The short cover letters accompanying the CVs

were also worded differently, while remaining standard. Each candidate was attributed a postal address, a mobile telephone number, and an e-mail address.

To ensure that the style or content of a particular application did not systematically influence companies' choices with respect to a particular candidate - in spite of the precautions taken in constructing them - we used a random permutation of the CVs between the identities of the fictitious candidates. The CVs and cover letters were thus alternated between the different candidates.

Applications for each job offer were sent on the day the offer was posted on the Internet, at an interval of several minutes, by e-mail from each candidate's mailbox. The response is considered positive when the recruiter invites a candidate to an interview or asks for further details about the candidate's current situation or qualifications. On the contrary, the response is considered negative if the recruiter ignores or formally rejects the application. When a candidate is invited to an interview, he/she declares that he/she already found a job and declines the offer.

3.2 Results

Since the candidates share a high level of qualification and a similar experience in a growing line of business, we did not expect to find a significant degree of hiring discrimination. Even if there were some discriminatory practices, they should be restrained by the difficulty in finding workers. That, however, is not at all what we find; discrimination clearly drives the data.

First, discrimination by origin is the most important type both by the range of candidates it applies to and by its degree. Second, gender discrimination is present for African origin women but is clearly weak or absent for women of French and Vietnamese origin.

We present the standard discrimination coefficients first and show how accounting for strong discrimination (FOSD) puts some light on these first-pass results. This will clearly show the contribution of this new measure to the analysis of hiring discrimination.

Table 2 presents the gender standard discrimination coefficients separately for each origin. We compare the success rate of each woman to the success rate of the man with the same origin. A positive coefficient indicates discrimination against women, while a negative coefficient indicates discrimination against men. We find that women with a French or African (Senegal, Morocco) origin are discriminated against, while men with a Vietnamese origin are discriminated against women of the same origin. The magnitudes of the discrimination coefficients are not comparable, however, since they use a different reference (the man of same origin varies from one coefficient to another).

For instance French-origin women have a discrimination coefficient of 5,2%, and Senegalese-origin women have a coefficient of 5,5%, but they do not face the same situation at all. If we use the strong discrimination criterion, we find that French-origin women are not FOSD by French-origin men (Table 5A), while Senegalese-origin women are FOSD by Senegalese-origin men (Table 5C). This is one case in which the new measure is useful. The two remaining cases (Moroccan and Vietnamese origins) fit with the FOSD criterion, so that Moroccan-origin women (Table 5B) and Vietnamese-origin men face strong discrimination (Table 5C). We can therefore propose the following summary: French-origin women face weak gender discrimination, while African-origin women and Vietnamese-origin men face strong gender discrimination. We will show later that it is possible to provide a full ranking of the candidates according to the weak and strong discrimination criteria.

For now, let $A \succ_s B$ denote a strong preference for candidate A (a strong discrimination against candidate B), and $A \succ_w B$ denote a weak preference for candidate A (a weak discrimination against candidate B). We have found that:

French-origin men \succ_w French-origin women
 Moroccan-origin men \succ_s Moroccan-origin women
 Senegalese-origin men \succ_s Senegalese-origin women
 Vietnamese-origin women \succ_s Vietnamese-origin men

One interesting property of the strong and weak discrimination criteria is that they allow ranking all the candidates of the testing along a single discrimination ordinal scale.

Table 3 presents the origin discrimination coefficients separately for each gender. Whether we consider women or men, one result never changes: foreign-origin candidates are always discriminated against to the benefit of the French-origin candidates. But to obtain this result, we need the FOSD criterion. Indeed, a close look at the discrimination coefficients reveals that the coefficient of Vietnamese-origin women is positive (2,6%) but not significant, so that, according to this coefficient, there would be not discrimination against Vietnamese-origin women. But a look at Table 5A shows that French-origin women FOSD Vietnamese-origin women, so that there is strong discrimination against Vietnamese-origin women. If we compare all of the candidates according to the strong discrimination criterion, we obtain a partial ordering:

(French-origin men, French-origin women) \succ_s (Moroccan-origin men, Vietnamese-origin women)
 \succ_s (Senegalese-origin men, Vietnamese-origin men)

🇵🇸 Moroccan-origin women

🇵🇸 Senegalese-origin women

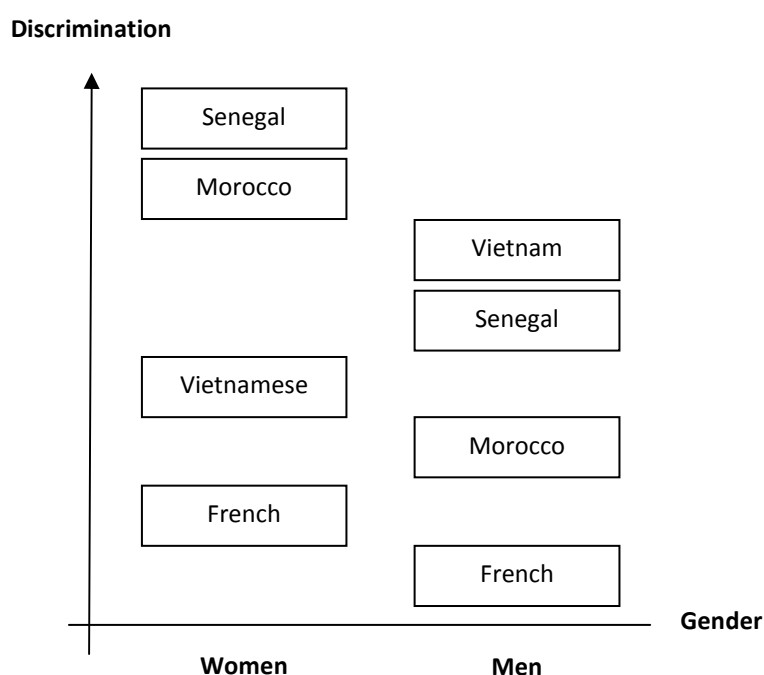
This partial ordering provided by FOSD can be completed by applying the weak discrimination criterion. We have already seen that French-origin men are preferred to French-origin women. To complete the ordering, we just need to compare Moroccan-origin men to Vietnamese-origin women, and Senegalese-origin men to Vietnamese-origin men. This can be done by using the additive property of weak discrimination: $\Delta_1(A, C) = \Delta_1(A, B) + \Delta_1(B, C)$. Using this property and the information contained in Tables 2 and 3, we obtain the complete ranking of the candidates by the recruiters:

French-origin men 🇵🇸 French-origin women
🇵🇸 Moroccan-origin men
🇵🇸 Vietnamese-origin women
🇵🇸 Senegalese-origin men
🇵🇸 Vietnamese-origin men
🇵🇸 Moroccan-origin women
🇵🇸 Senegalese-origin women

This ranking provides a position of the candidates along an ordinal discrimination scale, which is illustrated by Figure 4. In order to simplify the reading, we have separated genders. We obtain three main results. First, the candidates can clearly be separated into several groups. Among men, those of French and Moroccan origins face the weakest discrimination, while those of Vietnamese and Senegalese origins face a strong degree of discrimination from recruiters. Second, among women, those of French and Vietnamese origins are clearly preferred, while African-origin candidates face the strongest discrimination. Second, women are more often discriminated against than men, which is illustrated by an upward shift along the discrimination ordinal scale compared to men.

A potential critique that could apply to our methodology is the possible use of alphabetical ordering by recruiters as opposed to a discriminatory ordering. Therefore, we have checked whether alphabetical ordering had been used for each job offer and found that the data patterns are compatible with alphabetical ordering in only 11.5% of the cases (Table A1 in appendix 2). Therefore, the results that we present may reflect the recruiters' ranking rather than the neutral alphabetical ordering.

Figure 4: Discrimination ordinal scale



Conclusion

We have performed a correspondence testing in a line of business with a high job vacancy to applicant ratio. Considering young workers with a Master's degree and no period of unemployment, we did not expect to find evidence of significant discriminatory practices. The experiment indicates the contrary: overall, we do find evidence of significant discrimination by gender and origin. Moreover, in most cases, the degree of discrimination found is strong, since it corresponds to a pattern of stochastic dominance. While the standard discrimination coefficients go in the expected direction in most of the cases, they tend to underestimate the importance of discrimination in the French labor market. French-origin candidates of both genders clearly have more chances to be invited to a hiring interview than foreign-origin candidates, and African-origin women have the lowest chances to be invited.

Introducing strong discrimination (FOSD) brings value added to the analysis for two reasons. First, we find that candidates with similar standard discrimination coefficients do not face the same situation in the labor market, since some are weakly discriminated against and others are strongly

discriminated against. Second, candidates that do not seem to be discriminated against, because their discrimination coefficient is not significant although positive, can in fact face strong discrimination. Therefore combining strong and weak discrimination measures allows for a better assessment of the degree of discrimination that exists against the candidates.

REFERENCES

BERTRAND, M. and S. MULLAINATHAN (2003), Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination; NBER working paper series no. 9873; Cambridge: National Bureau of Economic Research.

DUGUET E. and P. PETIT (2005), "Hiring discrimination in the French financial sector: an econometric analysis on field experiment data", *Annales d'Économie et de Statistiques/Annals of Economics and Statistics*, n°78, pp 79-102.

FIRTH M. (1981), Racial discrimination in the British labor market. *Industrial and Labor Relations Review*, 34(2), 265-272.

FIRTH M. (1982), Sex Discrimination in Job Opportunities for Women, *Sex Roles*, 8 (8), pp. 891-901.

HECKMAN J.J. (1998), "Detecting Discrimination", *Journal of Economic Perspectives*, vol. 12, n° 2, pp 101-116.

KENNEY G. and D. WISSOKER (1994), "An Analysis of the Correlates of Discrimination Facing Young Hispanic Job-Seekers", *American Economic Review*, vol. 84, n°3, pp 674-683.

NEUMARK D., R. J. BANK, K. D. VAN NORT (1996), "Sex discrimination in restaurant Hiring: An Audit Study", *Quarterly Journal of Economics*, pp. 915-941.

RIACH, P. and J. RICH (2002), Field Experiments of Discrimination in the Market Place, *Economic Journal*, 112, pp. 480-518.

Table 2: Gender discrimination for each origin

The Student statistics have been computed by the block bootstrap with 10000 replications.
*: significant at the 5% level.

Origin :	France	Morocco	Senegal	Vietnam
No to both	63,5%	74,5%	83,2%	75,2%
Yes to man, no to woman	13,9%	15,2%	8,4%	4,8%
Yes to woman, no to man	8,7%	4,5%	2,9%	12,6%
Yes to both	13,9%	5,8%	5,5%	7,4%
Δ_1	5,2%*	10,6%*	5,5%*	-7,7%*
Student	1,96	4,33	2,93	3,30

Table 3: Origin discrimination for each gender

The Student statistics have been computed by the block bootstrap with 10000 replications.
*: significant at the 5% level.

Gender :	Men	Women
Origin : Morocco		
No to both	65,8%	72,9%
Yes to French origin, no to foreign origin	13,2%	16,8%
Yes to foreign origin, no to French origin	6,5%	4,5%
Yes to both	14,5%	5,8%
Δ_1	6,8%*	12,3%*
Student	2,70	4,90
Origin : Senegal		
No to both	67,7%	72,6%
Yes to French origin, no to foreign origin	18,4%	19,0%
Yes to foreign origin, no to French origin	4,5%	4,8%
Yes to both	9,4%	3,5%
Δ_1	13,9%*	14,2%*
Student	5,37	5,31
Origin : Vietnam		
No to both	68,7%	65,8%
Yes to French origin, no to foreign origin	19,0%	14,2%
Yes to foreign origin, no to French origin	3,5%	11,6%
Yes to both	8,7%	8,4%
Δ_1	15,5%*	2,6%
Student	6,06	0,90

Table 4: CDF of candidates' ranking













The table gives $\Pr(r_i \leq r)$, with $r \in \{1, \dots, 9\}$.

Gender	Woman	Man	Woman	Man	Woman	Man	Woman	Man
Origin	France	France	Morocco	Morocco	Senegal	Senegal	Vietnam	Vietnam
Rank	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
1	0,158	0,094	0,013	0,058	0,010	0,026	0,061	0,029
2	0,200	0,165	0,042	0,090	0,032	0,090	0,103	0,058
3	0,213	0,203	0,065	0,142	0,045	0,110	0,142	0,074
4	0,219	0,219	0,084	0,158	0,058	0,116	0,184	0,097
5	0,219	0,239	0,100	0,187	0,065	0,135	0,197	0,103
6	0,223	0,261	0,103	0,197	0,071	0,135	0,197	0,110
7	0,226	0,274	0,103	0,210	0,077	0,139	0,197	0,116
8	0,226	0,277	0,103	0,210	0,084	0,139	0,200	0,123
9	1	1	1	1	1	1	1	1

Table 5A: Stochastic dominance analysis

The table presents the differences between the CDFs of the candidates given in Table 4. Negative numbers have been underlined.









Reading: French origin candidates dominate all the other candidates, and do not dominate each other.

Reference :	Woman, France (1)							Man, France (2)					
Compared to :	Man France (2)	Woman Morocco (3)	Man Morocco (4)	Woman Senegal (5)	Man Senegal (6)	Woman Vietnam (7)	Man Vietnam (8)	Woman Morocco (3)	Man Morocco (4)	Woman Senegal (5)	Man Senegal (6)	Woman Vietnam (7)	Man Vietnam (8)
Rank	(1)-(2)	(1)-(3)	(1)-(4)	(1)-(5)	(1)-(6)	(1)-(7)	(1)-(8)	(2)-(3)	(2)-(4)	(2)-(5)	(2)-(6)	(2)-(7)	(2)-(8)
1	0,065	0,145	0,100	0,148	0,132	0,097	0,129	0,081	0,035	0,084	0,068	0,032	0,065
2	0,035	0,158	0,110	0,168	0,110	0,097	0,142	0,123	0,074	0,132	0,074	0,061	0,106
3	0,010	0,148	0,071	0,168	0,103	0,071	0,139	0,139	0,061	0,158	0,094	0,061	0,129
4	0,000	0,135	0,061	0,161	0,103	0,035	0,123	0,135	0,061	0,161	0,103	0,035	0,123
5	<u>-0,019</u>	0,119	0,032	0,155	0,084	0,023	0,116	0,139	0,052	0,174	0,103	0,042	0,135
6	<u>-0,039</u>	0,119	0,026	0,152	0,087	0,026	0,113	0,158	0,065	0,190	0,126	0,065	0,152
7	<u>-0,048</u>	0,123	0,016	0,148	0,087	0,029	0,110	0,171	0,065	0,197	0,135	0,077	0,158
8	<u>-0,052</u>	0,123	0,016	0,142	0,087	0,026	0,103	0,174	0,068	0,194	0,139	0,077	0,155
Stochastic dominance	<u>No</u>	(1)  (3)	(1)  (4)	(1)  (5)	(1)  (6)	(1)  (7)	(1)  (8)	(2)  (3)	(2)  (4)	(2)  (5)	(2)  (6)	(2)  (7)	(2)  (8)

To be continued

Table 5B: Stochastic dominance analysis (continued from table 5A)






The table presents the differences between the CDFs of the candidates given in Table 4. Negative numbers have been underlined.

Reference :	Woman, Morocco (3)					Man, Morocco (4)			
Compared to :	Man Morocco (4)	Woman Senegal (5)	Man Senegal (6)	Woman Vietnam (7)	Man Vietnam (8)	Woman Senegal (5)	Man Senegal (6)	Woman Vietnam (7)	Man Vietnam (8)
Rank	(3)-(4)	(3)-(5)	(3)-(6)	(3)-(7)	(3)-(8)	(4)-(5)	(4)-(6)	(4)-(7)	(4)-(8)
1	<u>-0,045</u>	0,003	<u>-0,013</u>	<u>-0,048</u>	<u>-0,016</u>	0,048	0,032	<u>-0,003</u>	0,029
2	<u>-0,048</u>	0,010	<u>-0,048</u>	<u>-0,061</u>	<u>-0,016</u>	0,058	0,000	<u>-0,013</u>	0,032
3	<u>-0,077</u>	0,019	<u>-0,045</u>	<u>-0,077</u>	<u>-0,010</u>	0,097	0,032	0,000	0,068
4	<u>-0,074</u>	0,026	<u>-0,032</u>	<u>-0,100</u>	<u>-0,013</u>	0,100	0,042	<u>-0,026</u>	0,061
5	<u>-0,087</u>	0,035	<u>-0,035</u>	<u>-0,097</u>	<u>-0,003</u>	0,123	0,052	<u>-0,010</u>	0,084
6	<u>-0,094</u>	0,032	<u>-0,032</u>	<u>-0,094</u>	<u>-0,006</u>	0,126	0,061	0,000	0,087
7	<u>-0,106</u>	0,026	<u>-0,035</u>	<u>-0,094</u>	<u>-0,013</u>	0,132	0,071	0,013	0,094
8	<u>-0,106</u>	0,019	<u>-0,035</u>	<u>-0,097</u>	<u>-0,019</u>	0,126	0,071	0,010	0,087
Stochastic dominance	(4)  (3)	(3)  (5)	(6)  (3)	(7)  (3)	(8)  (3)	(4)  (5)	(4)  (6)	No	(4)  (8)

To be continued

Table 5C: Stochastic dominance analysis (continued from table 5B)

The table presents the differences between the CDFs of the candidates given in Table 4.

Reference :	Woman, Senegal (5)			Man, Senegal (6)		Woman, Vietnam (7)
Compared to :	Man Senegal (6)	Woman Vietnam (7)	Man Vietnam (8)	Woman Vietnam (7)	Man Vietnam (8)	Man Vietnam (8)
Rank	(5)-(6)	(5)-(7)	(5)-(8)	(6)-(7)	(6)-(8)	(7)-(8)
1	<u>-0,016</u>	<u>-0,052</u>	<u>-0,019</u>	<u>-0,035</u>	<u>-0,003</u>	0,032
2	<u>-0,058</u>	<u>-0,071</u>	<u>-0,026</u>	<u>-0,013</u>	0,032	0,045
3	<u>-0,065</u>	<u>-0,097</u>	<u>-0,029</u>	<u>-0,032</u>	0,035	0,068
4	<u>-0,058</u>	<u>-0,126</u>	<u>-0,039</u>	<u>-0,068</u>	0,019	0,087
5	<u>-0,071</u>	<u>-0,132</u>	<u>-0,039</u>	<u>-0,061</u>	0,032	0,094
6	<u>-0,065</u>	<u>-0,126</u>	<u>-0,039</u>	<u>-0,061</u>	0,026	0,087
7	<u>-0,061</u>	<u>-0,119</u>	<u>-0,039</u>	<u>-0,058</u>	0,023	0,081
8	<u>-0,055</u>	<u>-0,116</u>	<u>-0,039</u>	<u>-0,061</u>	0,016	0,077
Stochastic dominance	(6)  (5)	(7)  (5)	(8)  (5)	(7)  (6)	<u>No</u>	(7)  (8)

Appendix 1 – Experimental Design

The field of the study covers all offers of full-time jobs as a developer, with either fixed-term or permanent contracts, located in the Paris region. We tested all the job offers that came to our attention between the beginning of February 2009 and the beginning of April 2009. Overall, 310 job offers were tested, corresponding to the mailing of 2480 applications (8 x 310).

The 8 candidates differ only by their origin and gender. The CVs of all of the candidates were sent for each job offer. They indicated their age (25 years old), their nationality (French) and their family situation (single, no children) on their CVs. They followed a similar school and university curriculum: a scientific *baccalauréat* (A level), followed by a degree in computing, and finally a Master's degree in computing from one of the following universities of the Paris region ("Ile-de-France"): Paris Sud, Paris VI Pierre et Marie Curie, Paris VII Diderot, Paris VIII Vincennes Saint-Denis, Paris XIII, Versailles Saint-Quentin, and Marne la Vallée.

The descriptions of training courses and positions occupied since entering the labor market have been chosen to compensate for eventual differences in Master's degree specializations. Ultimately, their training and experience give the 8 candidates equivalent skills.

During the first and second years of their Masters' degrees, the fictitious candidates acquired several months of work experience. At the end of the second year, each of the 8 candidates was recruited by the company that had hosted them during their work experience. Since then, they have acquired two years of experience as designer/developer in those companies. They are applying for the same type of position, often involving team management.

They all declare the following computer skills in their CVs: Programming: C, C#, C++, Java, XML, SCILAB, PHP, .net, J2EE ; Environments: UNIX, LINUX, WINDOWS ; Web development: Ajax, Web.2, HTML, JavaScript, .NET, GWT, RAILS, SPIP ; Data bases: SQL-Server, TSQL, MySQL ; Project management: UML, MERISE, Rational Rose ; Protocols: TCP/IP, SSH, FTP.

Eventually, the CVs have been rotated.

Appendix 2 – Checking alphabetical ordering

We have investigated whether the ranking of the candidates on each job offer is compatible with alphabetical ordering. The results are presented in Table A1, and we find that this is not the case.

Table A1: Alphabetical ordering at the job offer level

Number of job offers	310	
Offers with at least one positive answer	148	47.7% of 310
Compatible with alphabetical ordering	17	11.5% of 148
Incompatible with alphabetical ordering	131	88.5% of 148

Note: The 8 candidates were sent to each job offer, so that the total number of observations is 2480.