Záróbeszámoló

OTKA PD 76234
A diszkriminációs gyakorlat kialakulása, fennmaradása és belülről fakadó letörése
2009. szeptember 1 – 2012. december 31.
Vezető kutató: Takács Károly

Rövid összefoglaló magyarul

A kutatásban néhány olyan alapvető mechanizmust modelleztünk, amelyek diszkriminációhoz vezetnek a munkaerő-felvétel során. Azt vizsgáltuk, hogy milyen strukturális feltételek valószínűsítik a diszkriminációt és milyen hálózati feltételek mellett kerülhető ez el leginkább. Ágens alapú szimuláció és kísérletek segítségével körvonalaztuk, hogy milyen hálózati és keretezési feltételek segítenek a diszkrimináció endogén csökkentésében. Ezekhez a célokhoz az EU 7. Keretprogram IEF Programjának segítségével nemzetközi kutatói együttműködés létesült F. Squazzonival, M. Castellanival (University of Brescia) és G. Bravoval (University of Torino). A laboratóriumi kísérleteket Bresciában és Budapesten végeztük. Kulcseredményünk, hogy nagy egyenlőtlenségek lehetnek a munkaerő-felvételben akkor is, ha nincs minőségbeli különbség a különböző kategóriákba tartozó munkások között és a munkaadók alapvető célja a legjobb munkások felvétele. Az ágens alapú modellünkben megmutattuk, hogy a társas kapcsolatoknak különböző hatása lehet a diszkriminációra attól függően, hogy milyen típusú kapcsolatokról beszélünk. Továbbá, mind a szimuláció, mind a laboratóriumi kísérletek megerősítették, hogy a túlzott elvárások magasabb diszkriminációs szintet eredményeznek. Az eredményeinket számos nemzetközi konferencián prezentáltuk, valamint nemzetközi folyóiratokhoz küldtük publikálásra. A kísérleti eredményekből egy bírált könyvfejezet a Wiley kiadónál és egy műhelytanulmány már megjelent Torinóban.

Rövid összefoglaló angolul

The Evolution, Maintenance and Endogenous Breakdown of Discriminative Practices

In this research project we modeled certain fundamental mechanisms that lead to discriminative practices in hiring. We analyzed under which structural configurations discrimination is likely to occur and under which network conditions they can be avoided. With the help of agent-based simulation and laboratory experiments, we demonstrated how discrimination can be decreased by social network and priming mechanisms. For these objectives and with the support of the FP7 IEF Program of the European Union, an international collaboration has been established with F. Squazzoni and M. Castellani (University of Brescia) and G. Bravo (University of Torino). Laboratory experiments have been conducted in Brescia and in Budapest. Our key finding was that large inequality in employment can be pervasive even when there are no differences in average quality between different worker categories and employers only strive for high quality workers. In our agent-based model we showed that social networks might have a different impact on discrimination depending on the type of network ties. In addition, both simulations and laboratory experiments confirmed that high aspirations can lead to a higher extent of discrimination. Results have been presented at several international conferences and are now submitted for publication to international journals. A reviewed book chapter by Wiley and a working paper in Torino summarizing experimental results have already been published.

Részletes beszámoló angolul

I. Motivation

The research project had its focus on the problem of discrimination at hiring decisions. Discrimination is a relevant societal problem that undermines equal opportunities and threatens social integrity and could lead to severe conflicts. There are different recognizable traits (gender, ethnicity, hair color) that often provide a basis for negative discrimination. These traits are irrelevant for actions like cooperation or for compliance to social norms. Still, people act differently towards persons with different traits, as if they were different with regard to cooperation motives or norm obedience. Conflicts might occur if statistical discrimination provokes prejudice, labeling, and stigmatizing that in practice mostly target only members of one, disadvantaged category.

In our work, we could successfully demonstrate that large discrepancies between employment rates of different groups can occur as unintended results of the hiring process even in the lack of existing differences between the groups. Take for instance as an example that men and women do not differ on average in a quality that is required for a job. We show that even if employers have no in-born biases in favor of men, men could be hired disproportionally.

In particular, we addressed a previously overlooked aspect of discrimination: how social networks mechanisms contribute to the establishment of discriminative practices. The important role of social networks is well documented for getting a job. It is not clear from the existing literature, however, how social networks could lead to less or more discrimination.

The role of social networks is important in job hiring for two different reasons. First, the affective content of relationships creates obligations and favors that make hiring and recommending friends likely. Humans have the natural tendency to take into account and care for the welfare of friends and acquaintances, thereby *disrupting* the basic logic of a „perfect" market. Second, network ties are important channels of *information* that can be used to decrease the information asymmetry present at hiring decisions. Information about qualities and opportunities travel via social networks, and in this way networks bring the market *closer* to perfection.

When social networks are used in hiring because of their affective content, they replace meritocratic processes in hiring and result in suboptimal allocations (Ioannides and Loury, 2004; Petersen, Saporta, and Seidel, 2000; Tassier and Menczer, 2008). In this perspective, employers opt for hiring friends and acquaintances because of their personal commitments and due to their easy availability. They act according to their social self rather than directly maximizing their market success. The extended use of informal job search methods is believed to have a negative effect on the rate of mobility from low status to high status jobs (McBrier, 2003: 1212). If one of the groups has a better access to informal job search, then this is detrimental for the other group, as in the case of referrals from the "old boy network" in a wide range of fields (Rogers, 2000; McBrier, 2003).

When referral networks are in use and they are highly segregated, for instance, by ethnic group membership, they cause labor market segregation (Model, 1993; Tilly, 1998;

Elliott, 2001). Disadvantaged groups, especially with language deficiencies, rely more likely on insider referrals than advantaged (majority) groups, which further downgrades their mobility chances (Elliott, 1999; 2001; Green, Tigges, and Diaz, 1999). Members of a particular ethnic group tend to recommend their friends with the same ethnic background. Therefore, the employment statuses of path-connected workers are correlated (Krauth, 2004; Calvó-Armengol and Jackson, 2004; 2007). In case of segregated referral networks, groups with different reservation wages will receive different wages and firms can induce such segregation and discriminate between groups to increase their profit (Barr, 2009).

We used innovative methodologies of agent-based simulation and laboratory experiments with students as fictive employers to model and explain discrimination mechanisms. With these innovative methodologies, we were able to demonstrate that discrimination can be immense even in the lack of statistical differences between the qualities of the members of different group. Furthermore, we derived and analyzed predictions about how and which social structures help to eliminate discriminative judgments and practices. The results of our analyses contribute largely to the scientific understanding of discrimination and have practical implications for social policy to successfully tackle discrimination.

II. Results summary

Although many economists believe in optimal matching of jobs and employers in the labor market, in reality, dramatic employment discrepancies and mismatches can be experienced. This research project has shown that suboptimal situations occur even if there are no objective differences between groups of employees and employers are not biased initially.

First, we demonstrated by using a simple computational agent-based model that a certain level of employer discrimination is an inevitable consequence of asymmetric information and of limited supply of skilled labor. Moreover, our experiments found a prevalence of high discrimination rates in all experimental conditions, despite of the lack of quality differences between the groups.

Second, our simulations illustrated how higher aspirations of neutral employers could lead to higher discrimination rates compared to the situation when employers have low aspirations. Employers with high aspirations were also more discriminative than employers with low aspirations in our experiments. This finding can enrich existing explanations of why we experience more employment discrepancies in high status jobs and how can employers be trapped in their determination.

Third, our agent-based simulations highlighted structural mechanisms that can be responsible for the maintenance of discriminative practices. We found that the use of referral networks in getting a job does not necessarily increase discrimination in employment. Hiring via social network contacts; which could either be worker referrals, friendship ties between employers and workers, or business contacts; lower discrimination rates compared to the market composed of isolated employers. While experiments supported some of our hypotheses regarding network effects, we have not received full support for our predictions in the laboratory. The extensive use of worker referrals, in particular, has increased discrimination, which was not in line with our simulation results, but fits probably better to urban legends. Moreover, experiments

highlighted that the interaction of business recommendations, worker referrals, and high aspirations that potentially best describes real situations is characterized with the highest rates of discrimination.
Fourth, our simulations revealed how assertive workers could get into high quality jobs even if assertiveness is not correlated with qualities and employers are not concerned with it. Hence, we could explain how a quality-wage mismatch at the labor market could be an unintended consequence of standard hiring practices of profit-seeking employers.

III. Agent based simulations

Summary

By building an agent-based model of the hiring process, we showed that social and economic suboptimal situations are the rule more than the exception and discrimination could take place even when there is no objective quality difference between groups of employees; employers do not have discriminative tastes; and employers do not learn biased practices from each other. Our simulation results showed that objective differences in employment are inevitable consequences of asymmetric information and limited supply of skilled labor. Secondly, our results indicated that higher aspirations of fair employers tend to generate higher labor market inequality. This can enrich existing explanations on why we see higher employment discrepancy in high status jobs. Thirdly, we showed that assertive workers end up in high quality jobs even if assertiveness is not correlated or even negatively correlated with quality.
With regard to structural mechanisms, we demonstrated that social networks might have a different impact on discrimination depending on the type of ties. While worker referrals had a U-shaped effect, with a few referrals reducing discrimination, the structure of business recommendations did not play a crucial role. We found that certain social network mechanisms might decrease discrimination compared with a market composed of atomized employers as networks can allow employers to overcome the sampling bias. We also found that ties between workers and employers that are characterized by their affective content decrease discrimination more than ties that are characterized purely by information exchange.

Model outline

In our model, we considered *job hiring decisions* for *fixed terms* in a stable labor market with a *fixed set of employers* complemented by a *fixed set of workers* (*applicants*). We considered one recognizable trait in the population of workers and two social categories of this trait with a fixed category membership ($N_1=N_2$). We assumed that there was a surplus of labor supply in the market. There were more workers than jobs (J) in each period, i.e., $N_1+N_2 > J$, but the surplus was not exaggerated.
We first assumed one-sided matching; where employers chose workers and workers automatically accepted any offer. We assumed that employers were perfectly neutral and did not belong to any of the social categories. Given that we were not interested in wage

discrimination, wage competition and in the behavior of the supply side of the labor market, we did not make any differentiations between jobs offered.

We assumed that workers varied in their quality, but the variation was independent of group membership. We explored two configurations. In the default configuration, we drew integer values for individual quality from a uniform random distribution in {0, 1, …, 19}. In the other configuration, individual quality could take any value and were drawn from a normal distribution with a mean of 9.5 and standard deviation of 3, irrespectively of group membership. We assumed that individual quality was fixed and did not increase with employment.

The goal of employers was to fill vacant positions in their firms. In each contract period (year), employers were asked to fill a predefined number of positions by hiring workers. For sake of simplicity, employers filled up available positions randomly. We assumed that employers did not have information about individual quality of workers until they hired them. After hiring, the worker's individual quality was remembered by the employer for a given period of time, indicated by the memory parameter *m*. This implied that employers have information on a larger pool in the first terms, but the pool size had an upper constraint such that employers did not keep record of the quality of workers whom they employed for a long time. This could be due to turnover in human resource management, destroying old records, or for any other reason. Limited memories prevented the system from becoming a market with full information.

Employers used different channels for hiring. We assumed that employers gave priority to workers in house who had a quality above their expected standards. We assumed that quality standards were fixed over time and did not vary among employers.

We distinguished two types of friendship relations that were relevant for hiring: friendship ties between employers and workers and friendship ties that connected workers in-house with other workers. We assumed that friendship created obligations between employers and workers. Friends received priority in employment, if they were not employed before. On the other hand, we assumed that friends were not re-employed if they did not meet the quality standards or their quality was below group reputation scores. It is important to note that we assumed fair employers who could have friends from both groups with equal probability. Employer-worker ties were fixed over time independently of work experience and worker quality.

Friendship ties between workers were important for hiring through worker referrals. In-house employees could recommend their unemployed friends to their employers for a job. Note that friendship was independent of quality. Therefore, we did not "hardwire" a tendency for homophily based on quality. On the other hand, we assumed that friendship ties between workers were strongly homophilous with regard to group membership. Just like in other cases, if the recommended worker turned out to be of low quality, the employer hired somebody else.

Moreover, employers could obtain true information on individual quality of previously hired workers of business partners and hired them if they were unemployed. We considered *business ties* as mutual relations that were *fixed* over time and did not imply any costs for the partners involved. Therefore, business ties were channels of information about individual quality, but depending on the strength of influence in network ties (captured by a "group gossip" parameter *g*), they also shaped group reputation. We assumed that each employer formed and updated *group reputation* in each contract term.

Group reputation was the aggregation of individual quality information the employer encountered weighted by the reputation information the employer received from business partners. Therefore, it was *unique to each employer* and calculated as a weighted mean of individual experience (average quality of previously hired workers from the given group) and social influence from business partners.
In a given contract period, employers were selected in a random sequence to fill one of their jobs. This procedure was repeated until vacancies were available. When an employer was selected, the hiring decision followed this procedure:

1. *Re-hiring*: The employer checked previous employees and re-hired the one with the highest quality without a new contract; if this quality met his standards and exceeded group reputation scores.
2. *Hiring friends*: In case the vacancy could not be filled, the employer hires an unemployed friend with unknown quality randomly.
3. *Business recommendations*: In case the job was still not filled, one unemployed worker, who had the highest quality from previously hired workers of business partners was selected; if this quality met the standards of the employer and was higher than group reputation scores.
4. *Worker referrals*: In case the job was still not filled, one unemployed friend of recently employed workers of unknown quality was hired a) randomly or b) in order of the quality of referents.
5. *Hiring from the market*: In case the job was still not filled; one unemployed worker of unknown quality was hired randomly from the group with higher reputation. If group reputation was equal, an unemployed worker with unknown quality was randomly selected.

Note that our model ignored many aspects that economists would consider central to any hiring decision. In particular, the model included no turnover cost and legal restrictions on firing, which existed in many countries. It also ignored differential wages. In our view, both of these factors would tend to lead employers to stick with employees who were good enough rather than the best ones. It is worth noting that our employers followed satisfying rather than optimizing criteria as they kept in house all workers who had a quality higher than a threshold determined by the employer standards and group reputation. Our agent based model has been implemented in NetLogo (Wilensky, 1999). Our model has also been re-implemented in Repast to verify implementation coherence and to check internal validity. No difference has been found between the results of the two implementations.

Results: one-sided matching

We first report results from the model in which there were no network effects and hiring was based only on information based on personal experience of employers. In this case, employers satisfied with the quality of their in-house workers, re-hired them, while unsatisfied employers hired new workers on the market. Although discrimination was not expected, results showed that even under these ideal-looking conditions, employers could

be *perfect discriminators*. Even more surprisingly, a large disadvantage of one of the groups occurred also at the macro level.

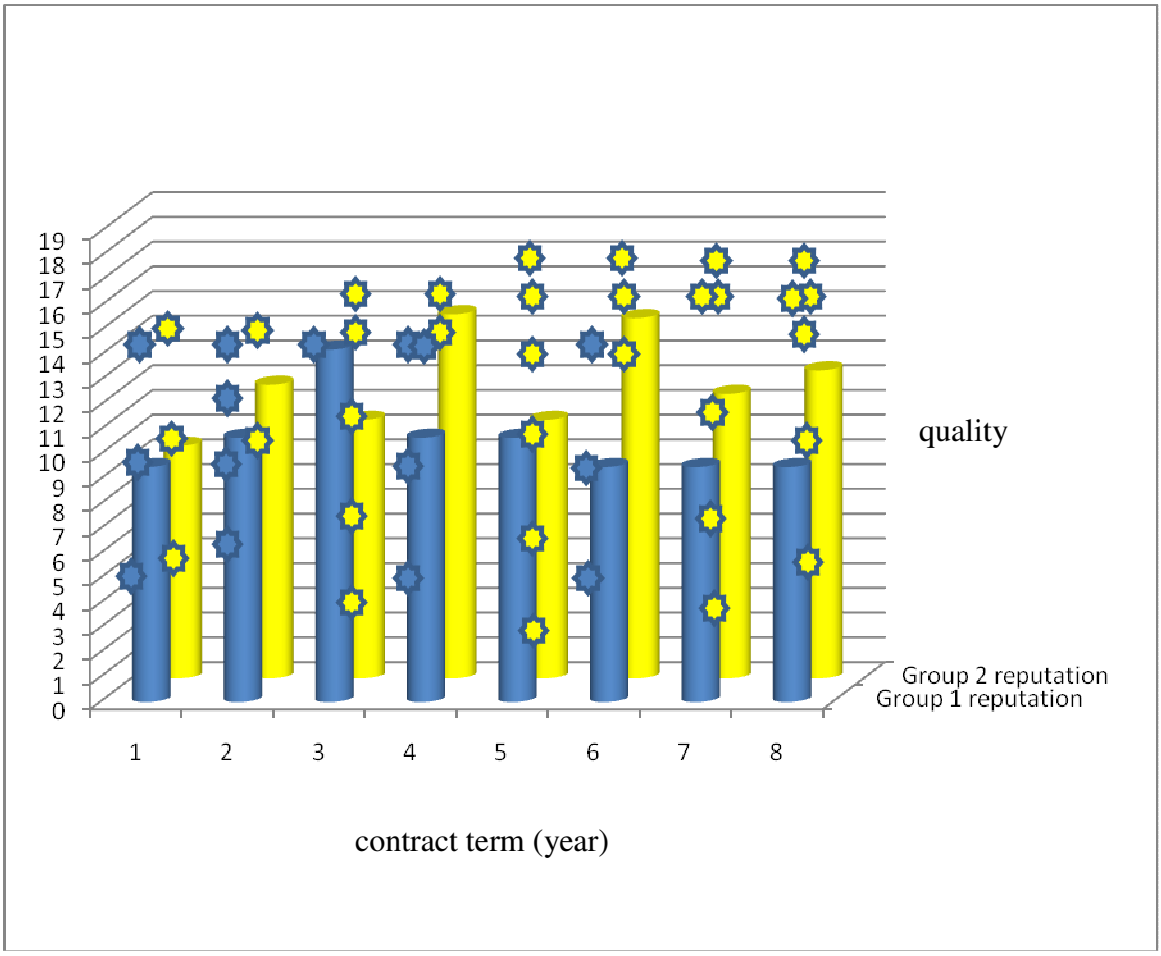

Figure 1. Expected composition of workers at a firm with one round of memory and 6 jobs.

Note: Colored bars indicated group reputation. Star symbols represented the expected quality of workers drawn from a random uniform distribution.

Figure 1 displays an illustration of a typical scenario of what could happen to a single employer. In the first hiring period, a small difference between the average qualities normally occurred due to random sampling. This subtle difference resulted in higher reputation for one of the groups, say the first one, and in the employment of new workers from this group. The expected average quality of these new workers, however, was below the average quality of workers from Group 2 who were kept in house in the second year. Therefore, the reputation score of Group 2 was higher and resulted in the hiring of four new employees from Group 2 in the third year. Yet again, the average quality of the new workers was below the reputation of Group 1, and therefore Year 4 was characterized by another change of fortune. By the time of the fifth contract term, employees we selected from one of the groups, only. This was once more counterbalanced in Year 6, but there was no way back from perfect discrimination and a solidified difference in group reputations from Year 7 onward. Hence, after quick switches in groups' fortunes, one

group gained an overall dominance, which remained stable over time. Therefore, perfect discrimination was the result of the rational search of fair employers.

Affective content: friendship between employers and workers

To measure discrimination in our simulations, we created a *macro level discrimination index* that looked at inequality in employment at the macro level. For equal group sizes and no differences in average quality, we defined the index as

$$\delta = \begin{cases} 1 - \dfrac{H_1}{H_2} & \text{if } H_1 \le H_2 \\ 1 - \dfrac{H_2}{H_1} & H_2 < H_1 \end{cases}, \qquad (1)$$

where $H_1 \le N_1$ was the number of hired workers from category 1. The index took the value of 0 when no discrimination took place and 1 when all jobs were filled with workers belonging to the same category.

If all employers were perfect discriminators and they perfectly discriminated groups randomly, then the δ index provided a low value. To correct for this problem, we created a *micro level index* δ_i that measured discrimination locally. The δ_i index simply compared the inequality in employment at the level of each individual employer and took the average of its distortion. It was calculated for each employer like δ, and then individual scores were averaged. The δ_i index took 0 if nobody discriminated and 1 if everyone discriminated perfectly one or the other groups.

The first major step in our model building strategy was to introduce friendship ties between employers and workers. We assumed that due to their affective content, these ties were unavoidable assets and burdens for the employers. They were assets in the sense that employers hired friends and could keep them committed if friends had high quality, while did not re-hire them if their quality turned out to be low. Friendship ties were burdens because unemployed friends with unknown working quality should receive priority at the hiring decisions.

It is important to note that at the set up of friendship ties between employers and workers, we did not assume *any bias*, i.e., employers had the same probability to have a link to workers in Groups 1 and 2. We manipulated the density of the bipartite friendship network between employers and workers and examined how the extent to which employers hired friends affected discrimination.

Our results showed that, under any parameter combination, discrimination did not increase with the increasing density of the bipartite friendship network. *Vice versa, if fair employers had and hired more friends, this helped to diminish discrimination* (see Figure 2). There was a drop in the discrimination indexes if employers had more friendship ties to employees. The drop was larger when employer standards were higher. For the higher employer standards, every employer was a perfect discriminator if they did not hire worker friends, and the discrimination rate was minimal if employers were friends with approximately 10% of the workers. In general, in case of a large density of employer-worker friendship ties, employers hired new workers only from their balanced friendship networks and not on the market.

Interestingly, hiring friends induced employers to behave more fairly compared with the situation where this type of social contacts did not exist, i.e., when employers relied on an extended and unbiased social network. In this case, their network pool was sufficiently large to ensure a suitable amount of workers with the required skills. To sum up, our results showed that commitment to friends, which was not driven by market incentives, efficiently reduced information asymmetry and consequently even discrimination.

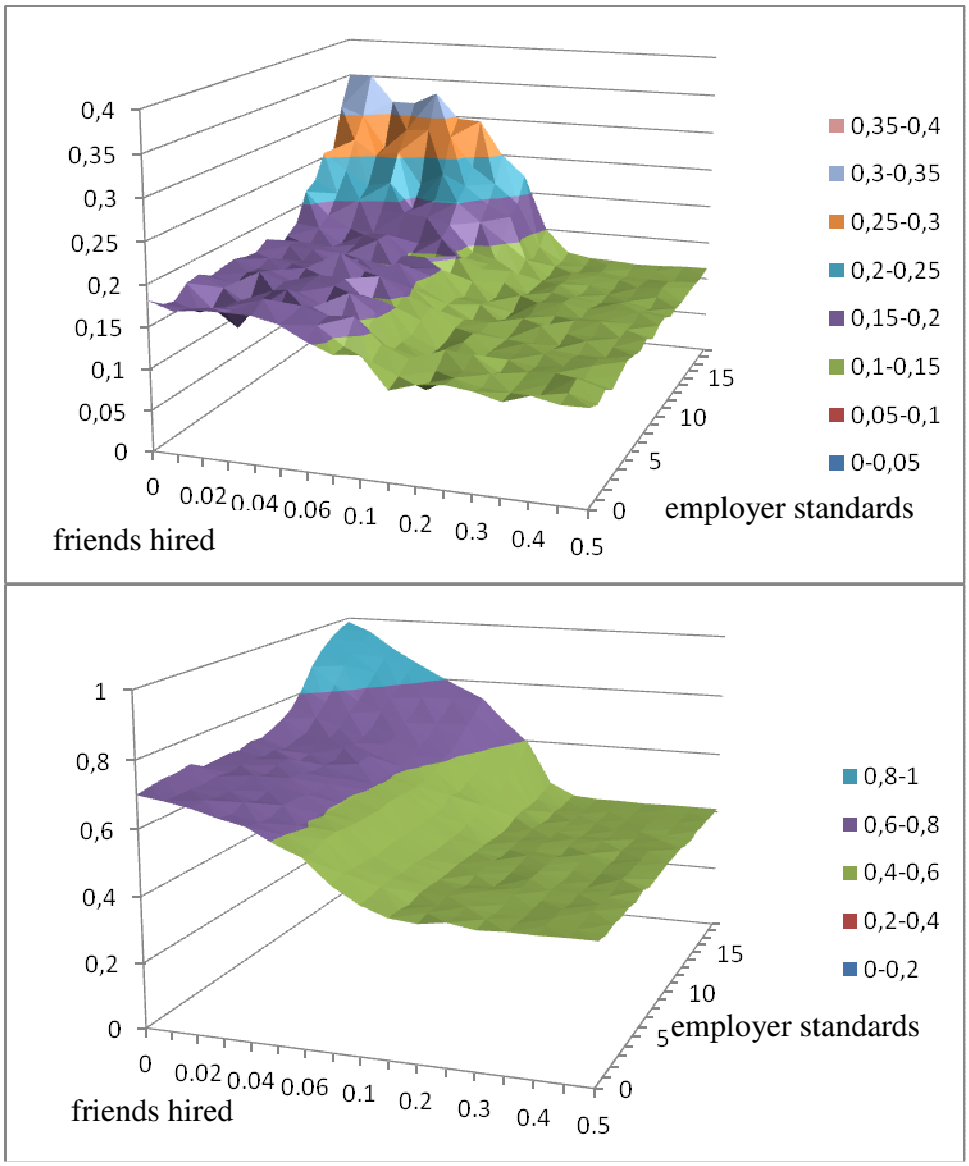

Figure 2. Mean values of the of the δ discrimination index (above) and of the δ_i micro level discrimination index (below) across 100 runs for each parameter combination; 34000 runs in total. Values are averaged for the density of the bipartite friendship network and for different employer standards.

Notes: 6 jobs per employer, no business contacts, m=5, no social ties among workers, 100 years per run.

Another interesting finding was that, independently of the density of the social network of workers, the density of the bipartite friendship ties had a strong effect on discrimination (Figure 3). Indeed, we found that the denser was the bipartite network, the lower the discrimination was. With no friendship ties between employers and workers, there was a stronger micro level discrimination, which decreased to a fair level thanks to hiring friends and worker referrals provided by friends hired before.

The social network of workers

Previous work brought us to believe that the increasing density of a segregated social network of workers could increase micro level discrimination. To test this, we manipulated the density of the *segregated* social network of workers. In the simulation setup of the social network, together with the required density, there were two perfectly segregated components of the social network among workers with an equal size. Within each component, ties were drawn randomly. Ties between the segments were created only above the critical value of density so that no more ties were possible within the components. During the hiring process, workers who were recently employed by an employer could "refer" their friends if jobs were open. Therefore, now employers could even pick friends of previously hired workers and benefit from the worker referral mechanism.

Results showed that discrimination quickly dropped with a *few* worker referrals (see Figure 3). This outcome was stable with or without business networks and in any kind of business networks. On the other hand, *many* worker referrals that characterized the denser worker network were detrimental for discrimination, especially in the lack of bipartite friendship ties.

This means that we found an interesting and robust U-shaped effect of social network density. Discrimination was higher in the complete lack and in the presence of sufficient ties among workers, and it was lower if only a couple of ties (40-100 ties in the numerical example of Figure 3 with 200 workers) were present. This U-shaped effect occurred in each network type, for all values of business network density and for all kinds of different specifications of the reputation mechanism.

It is worth noting that understanding the U-shaped effect is not easy. Results showed that a couple of ties between the workers (on average less than one) had an effect similar to the effect of employer-worker friendship ties. Indeed, the selection pool was enlarged and so discrimination dropped. More ties among workers gave rise to the expected effect induced by the segregation of contacts. This means that employers repeatedly hired from a dense circle of a single group of employees.

We found that labor market segregation did not increase by increasing the importance of referral hiring. On the other hand, labor market segregation showed a relatively stable and fair value across the parameter values of referral density and social network density, except for the lack of networks, when labor market segregation was higher. Our explanation is that business contact networks were the only information channel and alone they could not balance the outcome, as happened with referral networks.

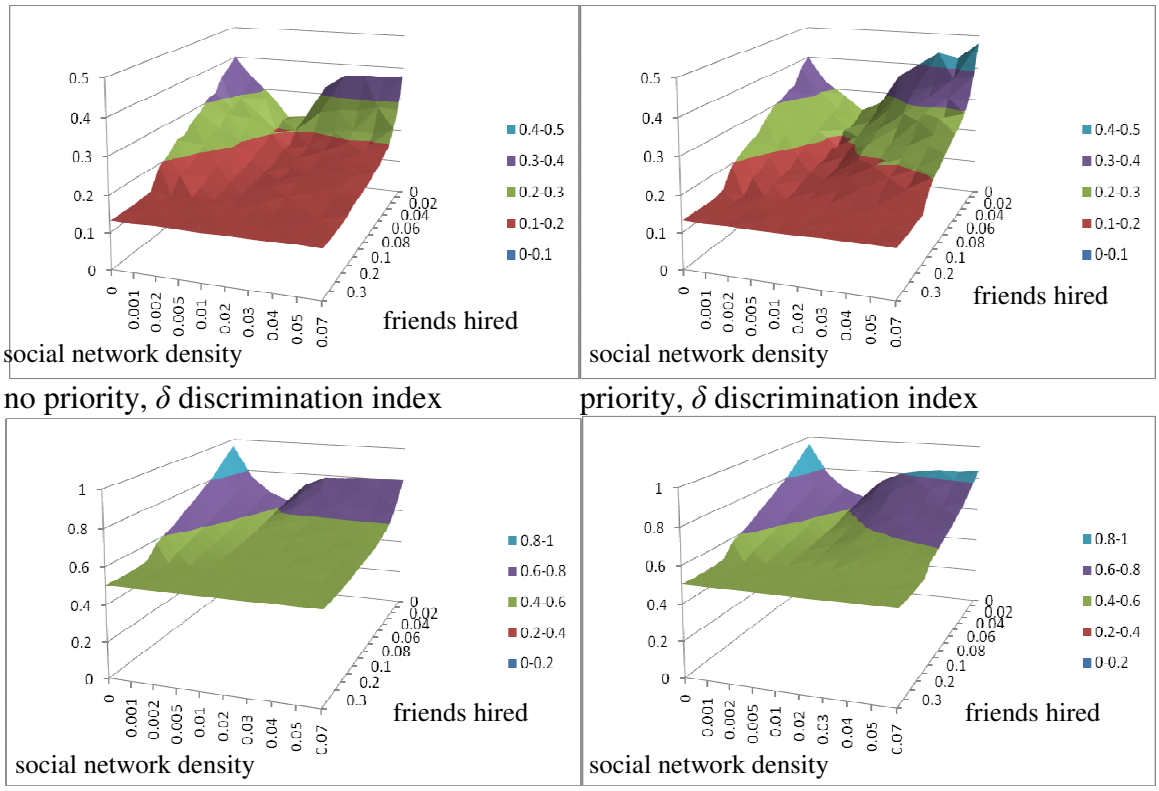

no priority, δ discrimination index | priority, δ discrimination index

no priority, δ_i discrimination index | priority, δ_i discrimination index

Figure 3. Mean values of the δ discrimination index (above) and the δ_i micro level discrimination index across 100 runs for each parameter combination (32000 for each part of the figure). Values are averaged for the density of social networks among workers and for the density of the bipartite friendship network between employers and workers. On the left: hiring in which no worker had a priority for referrals, on the right: hiring in which recommendations by workers with a higher quality were hired first.
Notes: 5 jobs per employer, random business network with a density of 0.1333; m=5, maximum employer standards (19); g=*0* and g=*0.1*, 100 years per run.

Furthermore, if worker ties were used for recommendations, the effect of employer standards on discrimination was diminishing. We found that employer standards had a stronger impact on discrimination rates both at the micro and at the macro level where social network ties among workers were not present.

Information exchange: business networks

As regards to the effect of network ties between employers on discrimination, we assumed that business contacts were fixed over time and did not imply any costs. As discussed before, business ties could be exploited to acquire important information about in-house worker quality. Therefore, these ties were as means to hire more skilled workers. Consequently, it was expected that the density and structure of business contacts could have an effect on discrimination. Our hypothesis was that higher density of business networks would have determined lower discrimination.

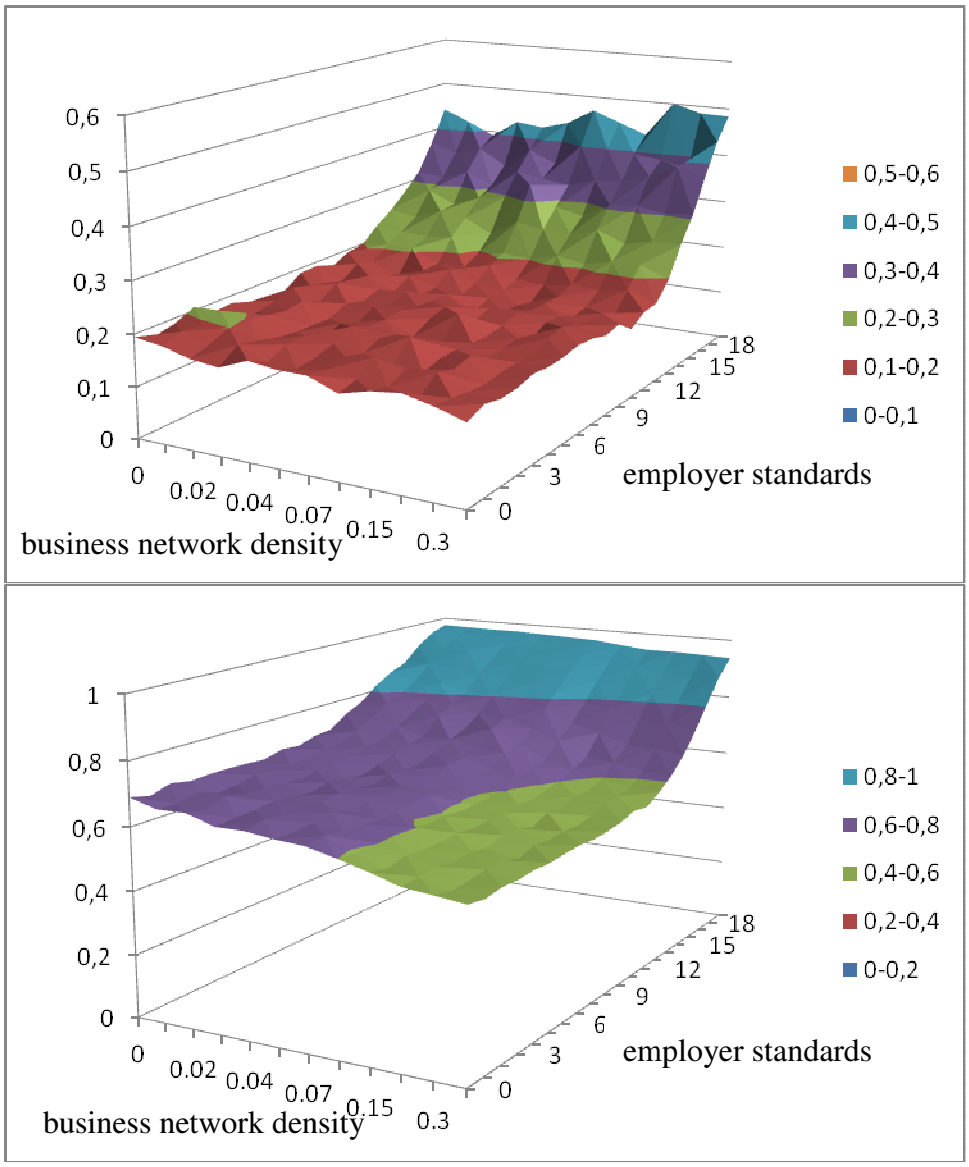

Figure 4. Average values of the of the δ discrimination index (above) and of the δ_i micro level discrimination index (below) across 100 runs for each parameter combination; 24000 runs in total. Values are averaged for the density of business networks among employers and for different employer standards.
Notes: 6 jobs per employer, random business network, no ties between employers and workers, m=30, 100 years per run.

As shown in Figure 4, we found that business network density decreased discrimination at the micro level, but only to a moderate extent. Small differences in the micro level discrimination index by the density of business contacts did not aggregate into significant differences in the macro level discrimination index. The effect of business networks was generally weak when compared with the impact of employer standards and worker referrals. Furthermore, business network density had a positive impact on the macro level discrimination index when employer standards increased. This small interaction effect

was present under a wide range of conditions, despite the fact that the micro level discrimination index decreased by business network density slightly even for higher employer standards.
These results indicated that the direct exchange of information about workers between employers could not have a strong effect on discrimination, thereby leaving the problem of information asymmetry in the labor market unsolved. Employers still showed bias, which were based on the overrepresentation of a few highly skilled in-house workers, especially in case of high standards.
Note that discrimination rates were low in the co-presence of business networks, employer-worker contacts, and worker social networks. Without worker referrals discrimination rates were higher, but still there was no difference between different business network structures. This means that between business partners, individual experiences took place more independently than expected and pairwise similarity of hiring choices between business partners was at a medium level. The choices of business partners were positively correlated, but only weakly. In contrast, without referral networks and exchange of group reputation, pairwise similarity in all types of business networks dropped to around zero, which indicated that simply merging available information on workers did not necessarily lead to correlated choices of employers.

Model extension: two-sided matching

In the two-sided extension of our model, we considered a job market, where employers selected employees and workers accepted or rejected offers. We assumed that workers had *assertiveness levels* of {0, 1, …, 19} taken from a uniform random distribution. On the other hand, all employers had the same quality standards (*aspiration levels*) of {0, 1, …, 19} or in other simulations, employer aspirations were also taken from a uniform random distribution. For the sake of simplicity, we assumed that assertiveness and aspiration levels were fixed over time.
The matching algorithm was the same as in the one-sided matching case, except that workers could reject offers. If a selected worker did not know the employer professionally (i.e., he/she did not work there before), he/she accepted the contract automatically. If the worker knew the employer (i.e., he/she worked there before), he/she rejected the proposal if the job quality was below his/her assertiveness level.

Results: two-sided matching

We found that higher aspiration levels of employers led to higher discrimination rates. This was a general result under a wide range of parameter values. Figure 5 shows that the aspiration trap occurred for aspiration levels higher than the mean of the scale. This means that there was no difference in discrimination rates of employers with low and medium level of aspiration. Rather, there was a linear increase in discrimination for aspiration levels higher than the mean of the scale.
For the highest aspiration level, Figure 5 shows that almost all employers were perfect discriminators. If they randomly favoured one or the other category, in the case of 15 employers, the expected value of the macro level discrimination index δ was 0.323. The fact that the observed value of δ was higher than the expected value indicated that there

was more discrepancy between the categories at the macro level than what we would have expected from the random distribution of individual discriminating tendencies. Note, however, that if the labor supply was radically extended (e.g., until ten times more workers than available jobs), then the aspiration trap faded away. The explanation is that, in this case, there was an abundance of workers with the highest quality and therefore all employers were satisfied easily with their selection.

Even if worker quality and assertiveness were perfectly independent and employers did not take assertiveness into account for selection, we found a concentration of most assertive workers in the best jobs. This result was obtained for a wide range of parameter values. In the example shown in Figure 5, the correlation between job quality and worker assertiveness was around one third for low aspiration levels of employers. The explanation is that while assertive workers accepted a job first, they easily stood up from their place if they were not satisfied, which brought success for them.

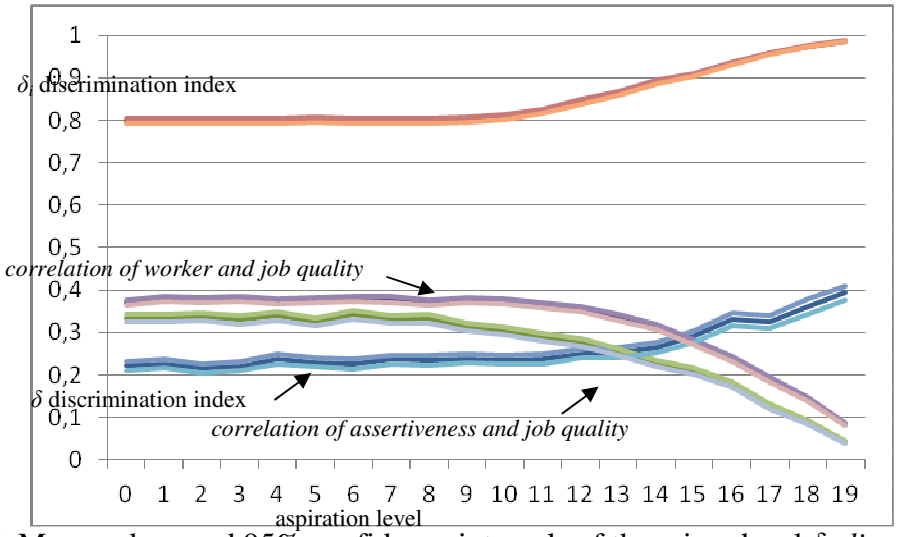

Figure 5. Mean values and 95% confidence intervals of the micro level δ_i *discrimination index* (*increasing curve above*), of the δ *discrimination index* (*increasing curve below*), of the *observed correlation between worker quality and job quality* (*decreasing curve above*), and of the *observed correlation between assertiveness and job quality* (*decreasing curve below*) in a total of 10000 runs when assertiveness and quality of workers were perfectly independent. Simulation parameters were as follows: 15 employers, 6 jobs per employer, N=200, m=10, each simulation run lasted for 100 contract years.

Although the assertiveness trap appeared under various conditions, it largely decreased if employers had the highest aspirations. Therefore, top jobs had the highest discrimination rates, but they were also more meritocratic in selection. In these jobs, the assertiveness of the worker did not play such a role because only the best workers were kept in house. The fact that assertiveness was less important to get the best jobs, however, did not mean that discrimination by employers could be justified.

An interesting finding was that higher aspirations even decreased the concentration of best workers in the best jobs. This means that desperate attempts to get the best workers not only increased discrimination and unequal handling of equal groups, but have also

resulted in a less meritocratic system, in which the best workers will not end up in the best jobs offered. This was due to the fact that there were no sufficient high quality workers in the market and employer aspirations were not correlated with job quality.
Furthermore, we manipulated the correlation between worker quality and assertiveness. We expected that increasing this correlation would have strengthened the concentration of assertive workers into the best jobs. This was confirmed as we found a positive relationship (see Figure 6). When more assertive workers had higher quality, their concentration into the best jobs increased. On the other hand, except for extreme values, an assertiveness trap occurred even in case of negative correlations between assertiveness and quality. This meant that the most assertive workers received the best jobs even if their quality was on average *lower* than that of less assertive workers. Note that assertive workers received better jobs even though employers did not take assertiveness into account for selection.

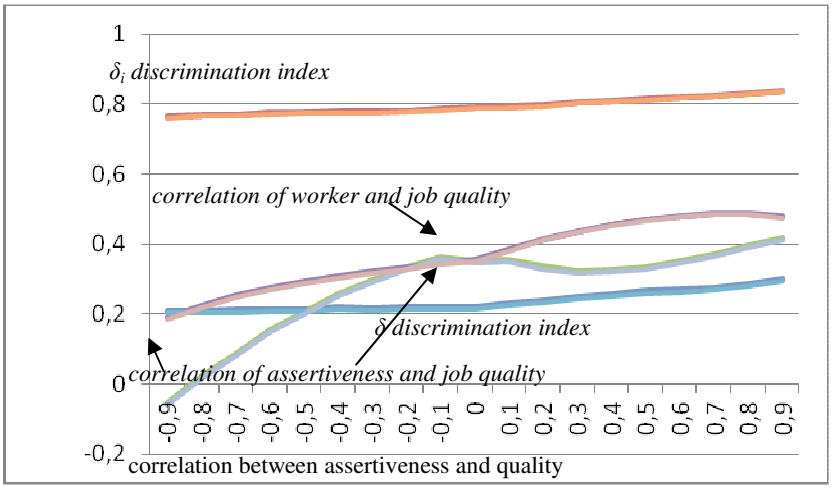

Figure 6. Mean values and 95% confidence intervals of the *micro level* δ_i *discrimination index* (*slowly increasing curve above*), of the δ *discrimination index* (*slowly increasing curve below*), of the *observed correlation between worker quality and job quality* (*quickly increasing curve above*), and of the *observed correlation between assertiveness and job quality* (*quickly increasing curve below*) in a total of 57000 runs. Simulation parameters were as follows: 15 employers, 6 jobs per employer, N=200, m={3, 5, 10, 20, 30, 50}, 0 aspiration level of employers, each simulation run lasted for 100 contract years.

Therefore, our results showed that the labor market never achieved a perfect matching of worker quality and job quality. The concentration of better workers in better jobs improved when the correlation between assertiveness and quality increased, but it did not imply a close match of worker and job quality. For instance, in the case of highest aspiration levels the correlation between worker and job quality never reached even the 0.14 level (Figure 7). The labor market mismatch between worker and job quality was larger for higher aspiration levels of employers (see Figures 7 and 8).
As discussed above, higher employer aspirations increased discrimination rates in general. A larger correlation between worker quality and assertiveness also increased discrimination rates, although less radically (see Figure 6). This was not true in case of

highest employer aspirations because micro level discrimination δ_i reached its theoretical maximum: every employer was a perfect discriminator (see Figure 7).

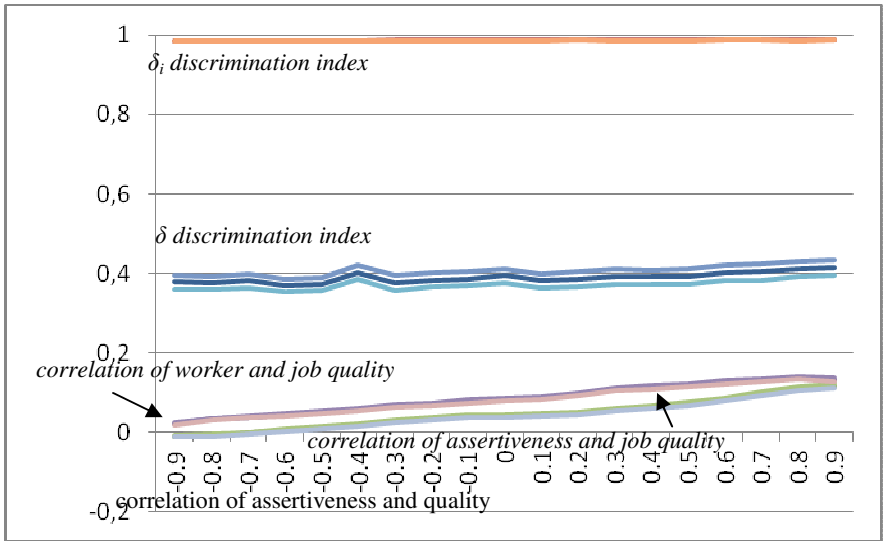

Figure 7. Mean values and 95% confidence intervals of the micro level δ_i *discrimination index* (*slowly increasing curve above*), of the δ *discrimination index* (*slowly increasing curve below*), the *observed correlation between assertiveness and job quality (quickly increasing curve below),* and the *observed correlation between worker quality and job quality (quickly increasing curve above)* in a total of 9500 runs. Simulation parameters were as follows: 15 employers, 6 jobs per employer, N=200, m=10, aspiration level of employers is 19, each simulation run lasted for 100 contract years.

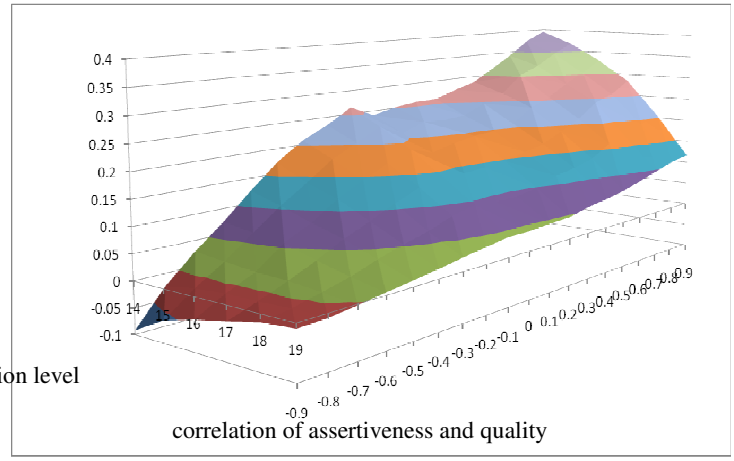

Figure 8. Average values of the *observed correlation of assertiveness and job quality* in a total of 57000 runs. Simulation parameters were as follows: 15 employers, 6 jobs per employer, 200 workers, m=10, each simulation run lasted for 100 contract years.

Figure 8 shows in detail the relationship between the correlation of assertiveness and quality and the occurrence of the assertiveness trap in the critical range of high aspiration levels of employers. It is possible to compare how the qualitative differences in the relationship in case of lowest and highest aspiration levels (Figures 6 and 7) were related to each other.

Next, we assumed that employers had heterogeneous aspiration levels. This was to check whether the aspiration trap and the assertiveness trap were simply due to the uniformity of employer aspiration levels. We also wanted to see whether the assignment of workers into jobs was more perfect under more realistic conditions, i.e., heterogeneous aspiration levels. Figure 9 shows that no qualitative difference was present if we relaxed the assumption of homogeneous aspiration levels.

Furthermore, we tested whether our results were robust against a wider range of parameter values. We have not found any major qualitative deviation from the above results. Among the control variables, memory of employers had a remarkable effect both on discrimination and on the correlation of assertiveness and job quality. This meant that longer memories enlarged the available pool of information and so decreased discrimination. On the other hand, micro-level discrimination did not extinguish even for perfect memories and longer memories have increased the extent of the assertiveness trap under any aspiration level and initial correlation of assertiveness and worker quality.

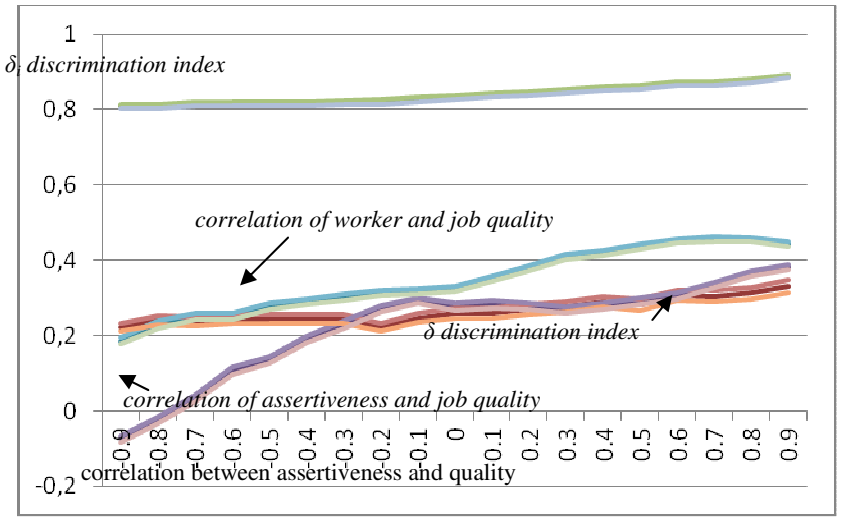

Figure 9. Mean values and 95% confidence intervals of the micro level δ_i *discrimination index* (*slowly increasing curve above*), of the δ *discrimination index* (*slowly increasing curve below*), of the *observed correlation between worker quality and job quality*, and of the *observed correlation between assertiveness and job quality* in a total of 9500 runs. *Notes*: 15 employers, 6 jobs per employer, N=200, m=10, aspiration levels of employers are taken from a random uniform distribution, each simulation run lasted for 100 contract years.

IV. Laboratory experiments

Summary

We tested two possible mechanisms that could influence discrimination in an artificial labor market, where subjects played the role of employers and were asked to choose workers from two groups labeled with colors. The first was a priming mechanism: we told subjects that there were sufficient high quality workers among both groups of

workers. The second was a network mechanism: subjects could benefit from true quality information about workers who had been previously employed by other "business contacts". Under each experimental condition, we found significant levels of discrimination. We found that subjects discriminated one group strongly more often when positive priming took place. The possibility of employing workers from business partners increased the average tendency to discriminate, despite the fact that this process reduced information asymmetry in the market. Results show that priming and social network mechanisms could have unexpected effects for employment inequalities and might increase discriminative bias. This would indicate that policy measures aimed at tackling discrimination should consider priming and network mechanisms with caution.
In another experiment, we checked whether referrals from other employers, referrals from workers, and higher quality standards requested by employers create inequality in employment in an artificial labor market. We found that referrals both from other employers and from employed workers increase discrimination in hiring, just as did the increase of employer standards. We also found that a combination of high quality standards and employers' referrals may decrease discrimination. Our results show that discrimination was higher when both types of referrals took place and employers had high standards

Method: Study 1

Experimental design

Our aim with the design was to test whether priming and network referrals could partly be responsible for higher or lower discrimination rates in an artificial labor market. We designed a simple environment, in which there was no any exogenous bias for or against any of the two groups.
Subjects were seated in a computer lab and listened to the instructions read by the experimenter. Twelve subjects were seated at a time in the laboratory and were randomly assigned to two groups of six. All interactions occurred anonymously through a computer network running the experimental software z-Tree (Fischbacher, 2007). Participants also received their instructions on paper and could read them on their screens while listening. After the instructions, subjects could openly ask questions. Later, no communication of any form was permitted and subjects received information solely from their screen. Subsequently a quiz tested whether subjects had understood the task correctly.
The experimental task was as follows. Subjects were asked to imagine that they were employers and were invited to hire workers for 8 vacant jobs in each of the 30 periods (a period was called a "year") forming the experiment. Each employment lasted for one period only. Subjects were paid at the end of the experiment according to the quality of workers they employed. More precisely, they earned their average period payoff, which was simply the total quality of hired workers divided by 8. Therefore, subjects were motivated to maximize the quality of their workers and to fill all their vacancies. In addition, subjects were warned by a computer dialog whenever they wanted to leave the stage without selecting the maximum number of workers. They could continue their selection after the dialog, if they choose to do so.

Each worker was an artificial agent characterized by an ID number, quality, and group membership. We used colors (blue and green) to differentiate between members of two groups of workers who were identical in size (60-60). Worker quality did not change during the experiment and subjects knew this from their instructions. Worker quality was assigned from a random uniform distribution in the range of {0, ..., 19}. Subjects were informed in advance about the range of qualities, but not about the means or distribution. As worker quality was assigned randomly, both groups had a mean quality of 9.5 and a uniform distribution of quality.

In each period of the game, subjects had to fill 5+3 vacancies with available workers simultaneously. To fill up the first five positions, they could select from: previously hired workers, workers of business friends (in the network condition) and random workers. Previously hired workers were displayed on the left side of the screen as a list, sorted by ID numbers. In each row, the ID number, group membership and their quality were displayed. This meant that subjects knew the quality of workers they had hired during previous periods and could hire them again. Subjects could select a worker by clicking on the appropriate row. Listed workers were those who were hired by the same employer in at least one of the two previous periods. Workers hired previous to these two periods disappeared, simulating memory limits or other real-world constraints.

The option of hiring random workers appeared on the right side of the screen in the form of three colored buttons with captions: hiring a random blue worker, hiring a random green worker, and hiring a random worker. Subjects were asked to hire workers within a limited time. The time remaining was displayed in the upper right corner of their screen.

The same worker could appear as an available option for multiple employers. In this case, the employer who picked the worker first could hire that worker. If a worker was hired before the decision limit, the corresponding row disappeared from every other list where the worker was displayed. On the screen of the subjects who had actually hired the worker, a new line appeared on the bottom of the screen indicating the ID and group membership of the worker and the number of vacancies left open. This made subject decisions in the experiment interdependent and put them somewhat in time pressure.

On a subsequent screen, subjects were called to select up to 3 new workers using the following options: hiring a random blue worker, hiring a random green worker, and hiring a random worker. This two-stage hiring process was intended to see whether discrimination occurred in general or only for new employment in our different experimental conditions.

At the end of the experiment and before being paid, participants were asked to complete a short questionnaire about their mood and motivations during the experiment. We measured computer use frequency as an indicator for disadvantages with regard to decision speed. Background questions and the participants' estimate about the average quality of workers were also asked. Earnings were paid in cash immediately at the end of the experiment.

Manipulations

We examined hiring decisions using a between-subjects 2 × 2 factorial design. First, we manipulated whether or not subjects were repeatedly primed with the information that

high quality workers were sufficiently available in both groups. The priming manipulation had two parts:

- The priming message was part of the experimental instructions.
- The priming message was re-iterated at the end of each period on a separate screen.

Secondly, we manipulated whether subjects could hire previous workers from their employer contacts. In the network condition, subjects were arranged in a circle network. We chose this setup because it is one of the simplest network structures, in which all positions are structurally equivalent. This meant that all subjects had two neighbors each. In the network condition, subjects could see the ID, the group membership and the true quality of workers who were hired in at least one of the two previous periods by their employer contacts. This list was displayed in the middle of their screen in the same way as their own workers were listed on the left side of their screen. These workers could also be selected by clicking the appropriate row. In the first period, lists were obviously empty. Note that as available options were largely overlapping in the network condition, if "quick" decision makers were discriminators, then the available pool for "slow" decision makers became biased in favor of the other group. This spillover effect, however, was compensated by a fairly longer list of available workers (own and of employer contacts) in the network condition.

Results of study 1

To analyze discrimination between treatments, we used various measures that all capture discrimination slightly differently. All measures, however, summarized one way or another the relative extent to which members of the two groups were employed. The first was an individual-level index $indDisc_{it}$ that was calculated as the relative difference between the numbers of workers hired from the two groups, considering colored choices only. So, a worker hired by a "hiring a random worker" option was excluded from the calculation. More precisely, $indDisc_{it}$ was calculated for subject i in period t as

$$indDisc_{it} = \frac{\left|H_{igt} - H_{ibt}\right|}{H_{igt} + H_{ibt}}, \qquad (1)$$

where H_{igt} was the number of knowingly green workers hired and H_{ibt} was the number of knowingly blue workers hired by subject i in period t. The index was zero when no discrimination took place and one when only workers belonging to the same group were hired.

The analysis of $indDisc_{it}$ indicated that discrimination was pervasive in our experimental labor market, despite the lack of average quality differences between the groups. The index value was significantly different above zero (t=77.322, p=0.000, one sided). We found a relatively high level of discrimination in all experimental conditions. The *lowest* value of $indDisc_{it}$ was 0.337 in the baseline condition, which corresponded to an average of 5.35 workers hired from one of the groups, and 2.65 workers hired from the other group, if all vacancies were filled with colored choices. Other experimental conditions showed even higher disparities.

In order to demonstrate the effectiveness of our manipulations, Figure 10 shows the average and dynamics of $indDisc_{it}$, by the presence of network manipulation (upper panel) and by priming manipulation (lower panel). It is interesting to note that discrimination increased dramatically after the first period in all experimental conditions. This indicates that subjects did *not* have strong prejudice or preferences for blue or green workers at the start of the experiment. However, discrimination *immediately increased* after subjects received their first feedback of worker quality and remained at this relatively high level throughout the experiment. As an ex post explanation, subjects used in principle *the first* information they received as an important and meaningful signal of group qualities and they quickly over-generalized to other potential candidates. The first experience therefore worked as an anchor and created a bias in one way or another.

Treatment	Baseline	Priming	Networks	Priming + Networks	Total
$indDisc_{it}$	.338	.347	.455	.442	.396
$strongDisc_{it}$	.26	.32	.20	.33	.28
$hiredtotalDisc_i$	.200	.194	.272	.251	.229
Decision N	906	1044	1044	870	3864
Subject N	30	36	36	30	132

Table 1. Discrimination measurements across treatments in study 1

Figure 10 also shows that values of $indDisc_{it}$ were higher in the network condition than in the no-network condition in all subsequent time periods. As there could be no network effects in the first period, data from the first period had to be excluded. Moreover, although the index provides a scale variable, due to integer values of H_{igt} and H_{ibt} and to the upper constraint of $8 \geq H_{igt}+H_{ibt}$, it can only take a few values. For this reason, we used non-parametric tests to compare index values between experimental conditions.

Table 1 shows the index values across experimental conditions. The network manipulation contributed to significantly higher values of $indDisc_{it}$ (Wilcoxon rank sum test W=3466745, p=0.000). The impact of network manipulation was also significant in terms of subject means of $indDisc_i$ with N=132 (W=3426, p=0.000). Considering the data from all periods, discrimination index values were not statistically different between the priming and no priming conditions (W=3654260, p=0.196). On the other hand, as suggested by Figure 10, the impact of priming manipulation may have decreased over time. This is reasonable, as subjects may be more influenced by the priming message at the beginning of the experiment, while potentially disregarding or not seriously considering the re-iterated message later. Looking at the values of $indDisc_{it}$ in the first 12 periods, differences were not significant between the priming and no priming conditions (W=618466, p=0.309, two sided). Hence, we can reject the hypotheses that priming had a significant effect on discrimination at the beginning of the experiment. Differences between the priming and the no priming condition in periods 13-30, however, were significant (W=1365457.5, p=0.005, two sided). The higher index values in the non-priming condition could possibly indicate that priming with positive messages may at least be effective in the long run in keeping back otherwise increasing discrimination tendencies. The decline of $indDisc_{it}$ in the priming condition, however, was very weak (from the 7^{th} period assuming a linear tendency: α=.430, β_t=-0.002 (t=-1.672, p=.095, not significant)).

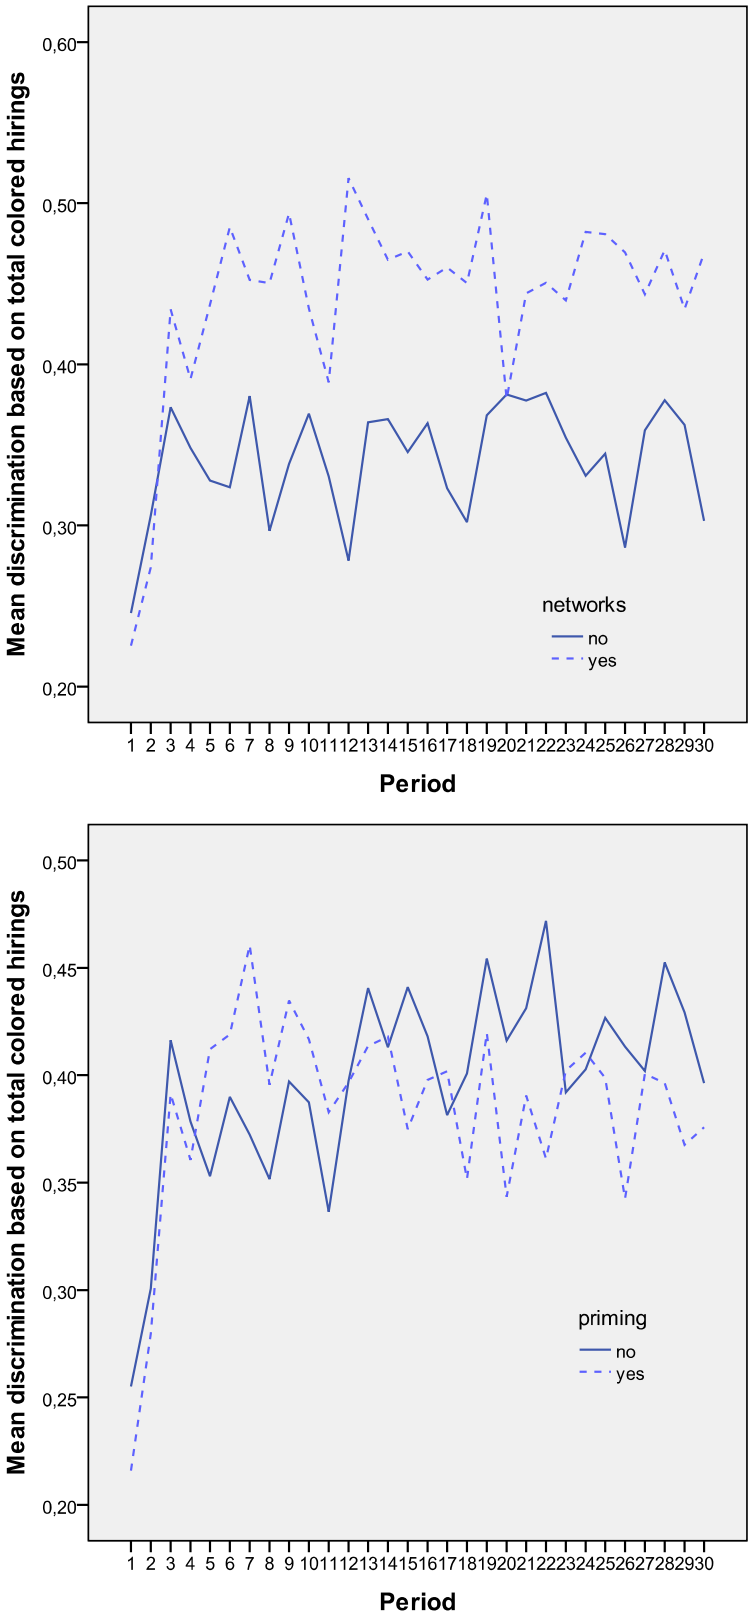

Figure 10. Mean values and dynamics of $indDisc_{it}$ by network manipulation (upper panel) and by priming manipulation (lower panel).

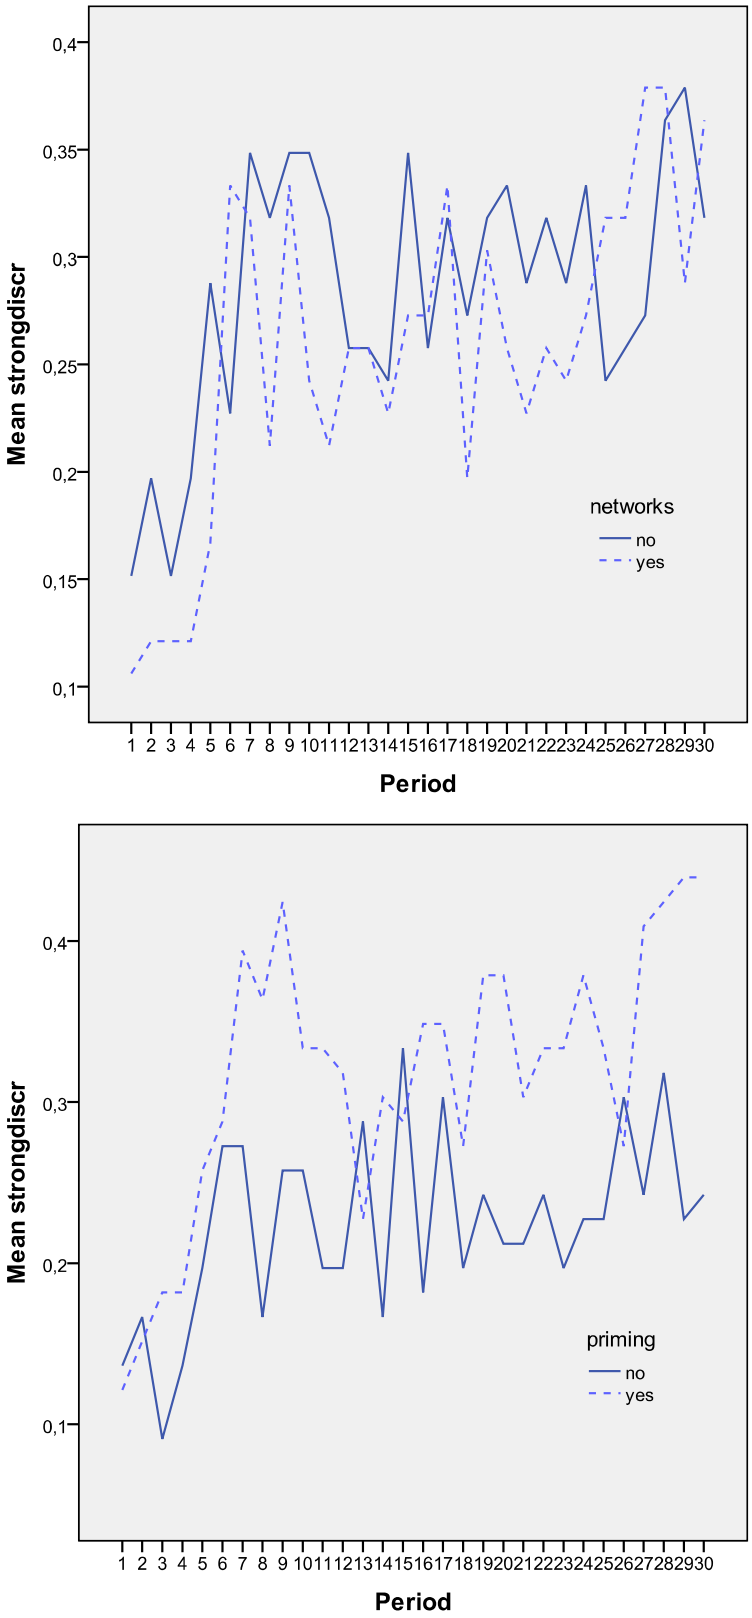

Figure 11. Mean values and dynamics of *strongDisc*$_{it}$ by network manipulation (upper panel) and by priming manipulation (lower panel)

The second measure we used was a binary variable $strongDisc_{it}$ which indicated whether subject i was a strong discriminator in period t or not. Strong discrimination meant that the subject hired only blues or only greens in the second stage of the task, when it was required to select up to 3 new workers.
Table 2 shows mean values of $strongDisc_{it}$, while Figure 11 provides a dynamic portrait. First, it is interesting to see that there were strong discriminators in every experimental condition. Second, it is important to note that this index was *not* aligned with values of $indDisc_{it}$ across the experimental conditions. Results showed that the network manipulation determined different and *lower* values of $strongDisc_{it}$ (W=3645636, p=0.048, two sided). Priming significantly *increased* the presence of strong discrimination (W=3587772, p=0.000). The impact of priming was also significant if we look at the subject means of $strongDisc_i$ with N=132 (W=3855.5, p=0.015, two sided). Figure 11 shows that strong discrimination also increased dramatically at the beginning of the experiment in all experimental conditions, but this upward shift took longer (7 periods) than the sudden increase in $indDisc_{it}$.
These results allow us to conclude that priming increased strong discrimination, but did not increase the average discrepancy between groups. Priming was probably effective for certain subjects but did not have a general effect. A second conclusion is that networks enlarged the impact of bias in general but decreased strong discrimination.

Experimental design of study 2

Subjects played anonymously in groups of six through a computer network, with two groups playing simultaneously in the same room to avoid identification of other group members. Subjects were asked to imagine that they were employers and were invited to hire ten workers per period, which represented one contract year. Workers were virtual agents and had ID numbers running from 1 to 200. Each worker had a fixed quality q_i drawn from a uniform integer distribution in the {0, ..., 19}interval at the beginning of the game. Half of the workers were labeled as "blues", half as "greens". Given that the color held no relation with the quality of workers, the distribution of quality was, on average, the same for blues and greens. This meant that any observed discrimination depended on the employers' evaluations and not on actual differences between workers' groups.
Subjects chose their workers one by one by selecting one of the following options: (i) hiring a worker randomly; (ii) hiring a blue worker randomly; (iii) hiring a green worker randomly; (iv) hiring one of the workers they hired in the previous period. To allow participants to use (iv), a list of the workers hired in the previous period, including their qualities and colors, was displayed on the screen. After the hiring stage (i.e., at the end of each period), participants were asked to provide an estimation of the average quality of both blues and greens.
In order to simulate turnover, retirement or worker mobility, in each round, workers had a 10% probability to be excluded from the list displayed to subjects. This meant that these workers became unavailable for hiring with the fourth procedure.
Each experiment lasted for 25 periods. Employers' profit depended on the quality of the hired workers, which was unknown before hiring them except for those hired in the previous period. To implement low and high standards experimentally, only the quality of

workers exceeding a given threshold τ was used for payment. The threshold was $\tau = 12$ in the low standard (–S) condition and $\tau = 17$ in the high standard (+S) condition. More specifically, profit was calculated by adding the qualities of workers with $q_i \geq \tau$ and by dividing the result by ten. Therefore, standards were exogenous constraints on decisions in our experiment.

At the end of each period, the average quality of blue and of green workers hired by the subject was calculated and displayed to the given participant along with the number of points earned. At the end of the experiment, profits were averaged across all 25 periods, with each point being exchanged with 1 Euro. At the end of the experiment and before being paid, participants were asked to complete a short questionnaire. Earnings were paid in cash immediately after the end of the experiment.

We examined hiring decisions using a between-subjects 2×2×2 factorial design. We manipulated these conditions: (i) whether high or low employer standards were imposed, (ii) whether employer referrals took place or not, and (iii) whether worker referrals took place or not (Table 2). Therefore, eight treatments, four testing for the pure effects of these factors and four testing for interaction between them, were contemplated (Table 2).

	Treatment	High standards	Employer referrals	Worker referrals
Pure effects	–S –E –W			
	+S –E –W	X		
	–S +E –W		X	
	–S –E +W			X
Interaction effects	+S +E –W	X	X	
	–S +E +W		X	X
	+S –E +W	X		X
	+S +E +W	X	X	X

Table 2: Treatment overview in study 2.

In case of employer referrals (+E), each participant was linked to one other "business friend" in a *directed circle network* with six nodes. In this condition, own workers hired in the previous period, and also those hired in the previous period by the business friend were displayed, including their color and quality. Furthermore, subjects received their friends' estimates about the average qualities of blues and greens. In the employer referrals condition (+E), subjects could also select a worker who was hired by the business friend in the previous period (selection procedure v). Note that this meant that the same worker could be hired by two or more participants in the same period and best workers were not subject to competition by business friends. Therefore, the decision speed was unimportant to determine who could hire the best workers. In case of no employer referrals (–E), no network existed and none of the subjects knew the workers hired by other subjects or others' estimations of workers' qualities.

In the worker referrals condition (+W), each worker hired in the previous period "recommended" a friend of the same color, randomly selected, without revealing its quality. The suggested workers were shown in a specific list on the participants' screen.

In case of worker referrals (+W), subjects could select a worker from this list (selection procedure vi). In case of no worker referrals (–W), participants did not have a list of recommended workers.

Results of study 2

Table 3 shows an overview of mean δ_{it} values in no-interaction treatments. All factor manipulations produced significantly higher discrimination than the baseline (Wilcoxon rank sum tests on individual averages: $W = 83$, $p = 0.006$ for +S; $W = 110$, $p = 0.052$ for +E; $W = 89$, $p = 0.010$ for +W; all p values are one tailed).

Treatment	N	Mean	SD	SE
–S –E –W	450	0.283	0.275	0.013
+S –E –W	450	0.426	0.301	0.014
–S +E –W	450	0.366	0.227	0.011
–S –E +W	450	0.438	0.313	0.015

Table 3: Average micro-level discrimination index in no-interaction treatments.

As our data included repeated observations of the same individual, a random effects (RE) model was performed. Table 4 shows the results of the model. The positive intercept indicates that a significant level of discrimination occurred in all treatments. This supported our first hypothesis.

	Estimate	SE	*t*	*p*
(Intercept)	0.257	0.032	8.051	0.000
+S	0.149	0.019	7.836	0.000
+E	0.090	0.020	4.584	0.000
+W	0.159	0.019	8.296	0.000
male	0.011	0.014	0.803	0.422
religion	–0.006	0.014	–0.422	0.673
economics	–0.001	0.016	–0.031	0.975
study year	0.011	0.005	2.186	0.029
part 2	–0.016	0.019	–0.868	0.386
$F(8,1791)$	11.759			0.000

Table 4: RE regression on the micro-level discrimination index in no-interaction treatments (*N=1800*).

Treatment	N	Mean	SD	SE
+S +E –W	450	0.224	0.206	0.010
–S +E +W	450	0.413	0.284	0.013
+S –E +W	450	0.477	0.334	0.016
+S +E +W	450	0.483	0.292	0.014

Table 5: Average micro-level discrimination index in treatments with interaction between factors.

As expected, worker referrals (+W) increased discrimination rates. On the other hand, employer referrals (+E) increased discrimination. Higher employer standards led to higher discrimination rates. It is important to note that all other variables included in the model, i.e., gender, professing a religion, studying economics (vs. all other faculties), the year of study, and a dummy included to check for differences between the first and second half of the game, had no significant effect except for a weak effect of the study year, with senior students who tended to discriminate more.

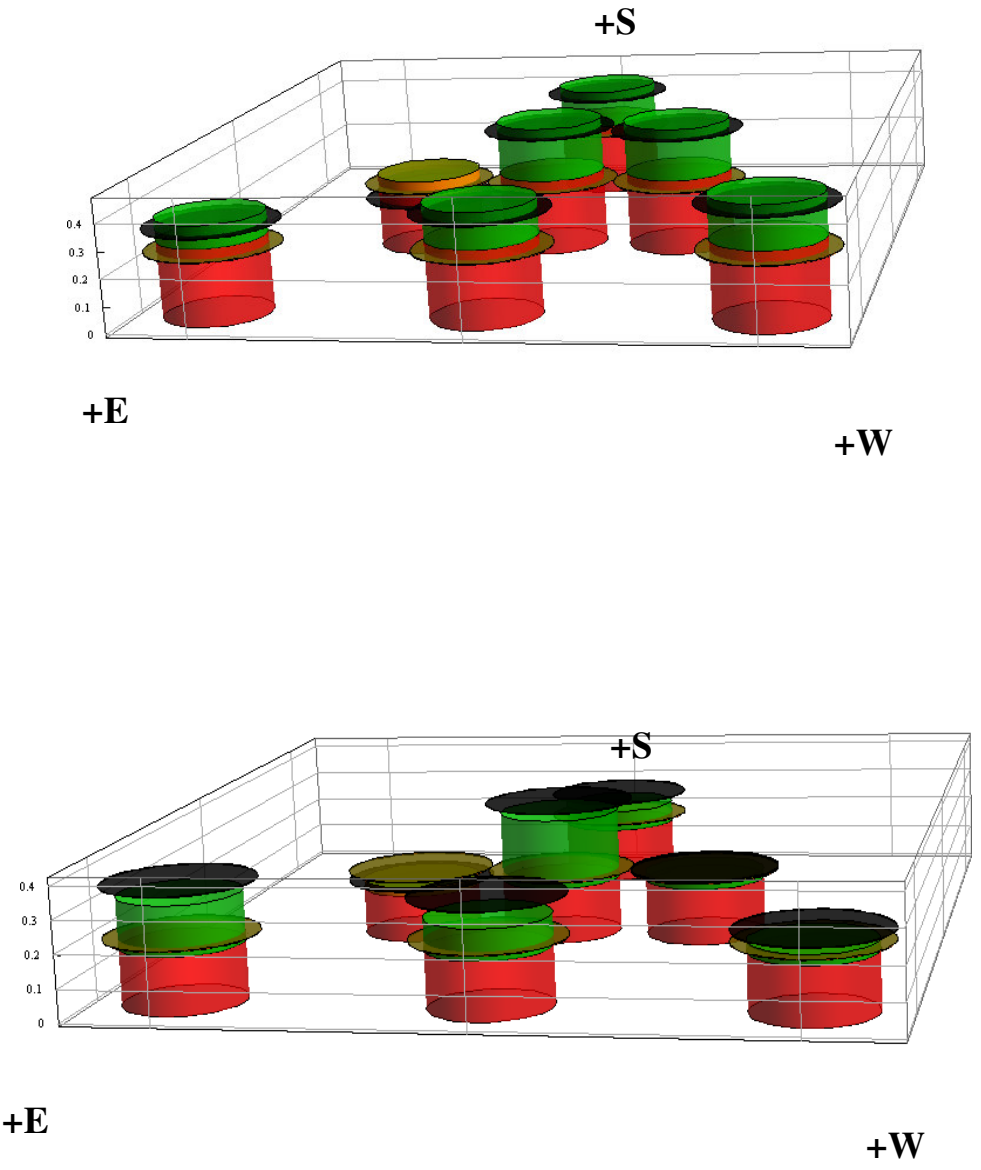

Figure 12. A comparison of model fitted predictions and average observed micro-level (upper panel) and macro-level (lower panel) discrimination across all treatments.

Treatments including interaction between the manipulated factors led to higher micro-level discrimination, except when employers with high standards used employer referrals (Table 5). All differences with the baseline (–S –E –W) were significant: –S +E +W, +S –E +W and +S +E +W led to higher discrimination ($W = 82$, $p = 0.005$, $W = 70$, $p = 0.002$, and $W = 43$, $p < 0.001$ respectively, all p values were one tailed) while +S +E –W led to lower discrimination ($W = 226$, $p = 0.022$, one tailed).

Figure 12 compares the results of treatments. Mean discrimination indexes from each treatment are represented with 3D columns. No-interaction treatments are displayed in the corners of the triangle, two-way interactions are represented on the edges and the three-way interaction is in the centre of the triangle. The –S –E –W baseline treatment is used as a basis for comparison. The red area indicates the observed value for the baseline treatment and the green area indicates the surplus in the observed values of the given treatment compared to the baseline treatment.

Results showed that the only exception was the +S +E –W treatment, in which observed mean values were below the baseline model. This difference is highlighted with orange. Fitted model predictions are displayed as black cylinders for all treatments. Orange cylinders indicate fitted model predictions for the baseline treatment as a basis of comparison. Black cylinders are the upper cylinders. Again, the only exception was the +S +E –W treatment, in which fitted model predictions were below the baseline model. A comparison of the upper and lower panel of Figure 12 confirms that there treatments had a different impact on the micro and macro discrimination indexes.

The RE model in Table 6, unlike the previous models, includes all treatments. All pure and interaction effects were significant in the new model. On the other hand, neither the individual variables nor the dummy, which indicated the second part of the game, were significant. The signs of all pure factor coefficients were positive, while binary interactions led to negative coefficients. On the other hand, only for +S * +E, the coefficient was sufficiently large to overcome the pure effect of the factors involved. Especially relevant was the large positive interaction effect of all three factors combined.

	Estimate	SE	*t*	*p*
(Intercept)	0.246	0.024	10.175	0.000
+S	0.148	0.019	7.822	0.000
+E	0.084	0.019	4.362	0.000
+W	0.161	0.019	8.476	0.000
+S*+E	–0.290	0.027	–10.714	0.000
+S*+W	–0.112	0.027	–4.152	0.000
+E*+W	–0.115	0.027	–4.257	0.000
+S*+E*+W	0.328	0.039	8.491	0.000
male	0.004	0.010	0.459	0.646
religion	0.014	0.010	1.389	0.165
economics	0.018	0.011	1.614	0.107
study year	0.004	0.004	0.951	0.342
part 2	0.000	0.014	–0.033	0.974
F(12,3587)	28.921			0.000

Table 6: RE regression on the micro-level discrimination index including interaction effects (*N=3600*).

V. Discussion

In reality, labor markets suffer from imperfections, inefficient outcomes, and a large mismatch between employer demands and worker skills. Although intrinsic qualities and skills are difficult to measure, there is hardly any perfectly meritocratic case in which the better workers would get the better jobs and no group would suffer from discrepancies in employment. Employment discrimination occurs largely because employers are unable to observe the individual qualities and skills of employees in advance and have to rely on signals and external characteristics at their hiring decisions. Certain recognizable traits are not the results of investments, but are the unalienable characteristics of individuals and as such, they could be used to differentiate between applicants. Gender, race and age are the most salient social categories immediately encoded in case of any interaction. When information is hard to collect, recognizable social category membership is often used to form group reputations and prejudices which are used as proxies to estimate and judge the individual abilities of category members. If there are statistical differences between the categories, we talk about statistical discrimination. The most difficult problems of discrimination, however, are those when discrimination is not based on statistical differences.

Our aim in this project was to examine why and how discrimination might emerge and be maintained in the lack of mean quality differences and considering neutral employers. In particular, we demonstrated that discrimination can be pervasive at a labor market also in the lack of mean quality differences and if employers with unbiased intentions strive for high quality workers. We have come to this conclusion using an agent-based simulation model that has been set up in NetLogo and re-implemented in Repast. Re-implementation has been rarely done before in social simulation studies, although such an exercise radically improves validity of the results.

We showed that high aspirations of employers increase discrimination rates for a wide range of parameters, which might explain why we observe larger inequalities in employment at top level jobs. Experimental results confirmed the aspiration effect and found a large extent of discrimination in general in all experimental conditions. We have run two series of experiments with different settings to arrive at valid conclusions. The main advantage of using experimental methodology compared to other empirical methods was that we could test the hypothesized associations and mechanisms unambiguously in a perfectly controlled environment.

Simulations highlighted that hiring via social networks, which could either mean employing friends, using worker referrals, or business recommendations, lowers discrimination rates compared to a market that is composed of isolated employers. These results have only received partial support in our laboratory experiments. In particular, in the experiments, more in line with the common belief, we found that referral hiring contributes to more discrimination at the labor market. Our future task is to get back to agent-based modeling and develop the model further to explain this discrepancy.

Practical benefits

Both understanding discrimination and fighting discrimination can benefit from our research that has investigated previously neglected aspects and underlying reasons of observed discrimination. Our research was primarily theoretical, and the innovative methods that have been applied were at a higher level of abstraction from reality. This had the major advantage, however, that we could analyze certain fundamental mechanisms with the perfect exclusion of rival explanations. Our findings about fundamental mechanisms can help to understand why anti-discrimination campaigns are not fully successful in practice, and can provide advice for forming efficient policies to fight discrimination and labor market segregation. Besides policy implications, we have something to say also for employers, business practitioners, and human resource managers about the fallacies of hiring decisions and best practices to avoid inefficiencies related to sampling biases.

References

Barr, T. 2009. "With Friends Like These: Endogenous Labor Market Segregation with Homogeneous, Nonprejudiced Agents." *American Journal of Economics & Sociology*, 68(3): 703-746.

Calvó-Armengol, Antoni, and Matthew O. Jackson. 2004. "The Effects of Social Networks on Employment and Inequality." *American Economic Review* 94: 426–454

Calvó-Armengol Antoni, and Matthew O. Jackson. 2007. "Networks in Labor Markets: Wage and Employment Dynamics and Inequality." *Journal of Economic Theory* 132: 27–46

Elliott, James R. 1999. "Social Isolation and Labor Market Insulation: Network and Neighborhood Effects on Less-Educated Urban Workers." *Sociological Quarterly* 40 (2): 199-216.

Elliott, James R. 2001. "Referral Hiring and Ethnically Homogeneous Jobs: How Prevalent is the Connection and for Whom?" *Social Science Research* 30: 401-425.

Fischbacher, U. 2007. "z-Tree: Zurich toolbox for ready-made economic experiments". *Experimental Economics*, 10, 171–178.

Green, Gary P., Leann M. Tigges, and Daniel Diaz. 1999. "Racial and Ethnic Differences in Job Search Strategies in Atlanta, Boston, and Los Angeles." *Social Science Quarterly* 80 (2): 263-278.

Ioannides, Yannis M., and Linda D. Loury. 2004. "Job Information Networks, Neighborhood Effects, and Inequality." *Journal of Economic Literature* 42: 1056–1093.

Krauth, B.V. 2004. "A Dynamic Model of Job Networking and Social Influences on Employment." *Journal of Economic Dynamics and Control*, 28:1185–1204

McBrier, Debra Branch. 2003. "Gender and Career Dynamics within a Segmented Professional Labor Market: The Case of Law Academia." *Social Forces* 81 (4): 120-1266.

Model, Suzanne. 1993. "The Ethnic Niche and the Structure of Opportunity: Immigrants and Minorities in New York City." In: Katz, M. (ed.): *The Underclass Debate: Views from History*, 161-193. Princeton, Princeton University Press.

Petersen, Trond, Ishak Saporta, and Marc-David L. Seidel. 2000. "Offering a Job: Meritocracy and Social Networks." *American Journal of Sociology* 106: 763-816.

Rogers, Jackie Krasas 2000. *Temps: The Many Faces of the Changing Workplace*. ILR Press.

Tassier Troy, and Filippo Menczer. 2008. "Social Network Structure, Segregation, and Equality in a Labor Market with Referral Hiring." *Journal of Economic Behavior and Organization* 66: 514–528.

Tilly, Charles. 1998. *Durable Inequality*. Berkeley: University of California Press.

Wilensky, U. (1999) *NetLogo*. http://ccl.northwestern.edu/netlogo/. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL.