

Filippin, Antonio

Working Paper

Discrimination and Workers' Expectations: Experimental Evidence

IZA Discussion Papers, No. 824

Provided in Cooperation with:

IZA – Institute of Labor Economics

Suggested Citation: Filippin, Antonio (2003) : Discrimination and Workers' Expectations: Experimental Evidence, IZA Discussion Papers, No. 824, Institute for the Study of Labor (IZA), Bonn

This Version is available at:

<https://hdl.handle.net/10419/20069>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

IZA DP No. 824

Discrimination and Workers' Expectations: Experimental Evidence

Antonio Filippin

July 2003

Discrimination and Workers' Expectations: Experimental Evidence

Antonio Filippin

*University of Milan, EUI
and IZA Bonn*

Discussion Paper No. 824
July 2003

IZA

P.O. Box 7240
D-53072 Bonn
Germany

Tel.: +49-228-3894-0
Fax: +49-228-3894-210
Email: iza@iza.org

This Discussion Paper is issued within the framework of IZA's research area *The Future of Labor*. Any opinions expressed here are those of the author(s) and not those of the institute. Research disseminated by IZA may include views on policy, but the institute itself takes no institutional policy positions.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent, nonprofit limited liability company (Gesellschaft mit beschränkter Haftung) supported by Deutsche Post World Net. The center is associated with the University of Bonn and offers a stimulating research environment through its research networks, research support, and visitors and doctoral programs. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public. The current research program deals with (1) mobility and flexibility of labor, (2) internationalization of labor markets, (3) welfare state and labor market, (4) labor markets in transition countries, (5) the future of labor, (6) evaluation of labor market policies and projects and (7) general labor economics.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available on the IZA website (www.iza.org) or directly from the author.

ABSTRACT

Discrimination and Workers' Expectations: Experimental Evidence*

This paper is an experimental analysis of the role played by workers' expectations in explaining the puzzling long-run persistence of observed discrimination against certain minorities in the labor market. The experiment provides some evidence supporting the theoretical prediction that unequal outcomes may emerge due to disadvantaged workers' wrong expectations of being discriminated against. However, this effect is not long-lasting, since players learn the true state of nature in later stages of the experiment, failing to generate a Self-Confirming Equilibrium driven by wrong beliefs. The strategy method provides additional evidence that expectations matter.

JEL Classification: C92, J71, J15, D84

Keywords: experiments, discrimination, workers' expectations

Antonio Filippin
Department of Economics
University of Milan
Via Mercalli, 23
20122 Milan
Italy
Email: antonio.filippin@unimi.it

* I would like to thank Pierpaolo Battigalli and Andrea Ichino for their excellent supervision, Thomas Bauer, Gary Charness and Karl Schlag for their valuable suggestions. Needless to say, all the remaining errors are mine.

1 Introduction

The goal of this paper is to provide experimental evidence concerning the role of workers' expectations as an explanation for the puzzling long-run persistence of observed discrimination against certain minorities in the labor market. The model used as a benchmark (see Filippin (2003)) shows that *ex ante* identical groups of workers may be characterized by unequal outcomes in equilibrium due to their different beliefs. In particular, the model shows that unequal outcomes may arise when minority workers wrongly believe that they are discriminated against, even when employers do not do so either directly or statistically.

The setting of the experiment is not a carbon copy of the model. First, the experiment is designed in a way that does not remind subjects that it deals with discrimination in the labor market. This is to keep their attitudes, experiences, and opinions on the matter out of the lab, preserving the importance of payoffs in driving their behavior. The participants are randomly assigned to two populations: red and blue. In every lottery there are only two participants, one from each population. In every trial every participant has an endowment of 10 Euro cents and bids in order to win a prize worth 25 Euro cents. The players know that they face only one opponent in every trial. They also know that the higher bid wins the prize, unless they face a "crazy computer" that awards the prize to the Red regardless of the bids. Bids are not given back to either the winner or the loser. Colors (red and blue) are the equivalent of gender (or race, etc.), the prize stands for the promotion and the "crazy computers" play the role of discriminatory employers. Finally, the amount bet plays the role of effort. In the model, workers can choose to give up some utility in the first period exerting an inefficiently high level of effort in order to increase the probability of being promoted. Similarly, in the experiment subjects can trade (part of) their endowment with a higher probability of winning the prize.

This paper analyzes whether common past experience works as a coordinating force. In other words, starting from a situation where there are discriminatory employers, do minority workers still expect of being discriminated against when they do not know that discriminatory tastes have disappeared? A positive answer to this question would provide evidence that historical factors are crucial in selecting one among different possible outcomes (path dependent equilibrium selection), pointing toward the existence of hysteresis. This would provide useful insights concerning the long-run persistence of discrimination in the labor market.

In the lab, a reduction of bids made by disadvantaged subjects leading to a lower fraction of prizes awarded to them, is observed in three out of seven sessions, in line with the theoretical predictions, but it vanishes rather quickly during the treatment, failing to generate a Self-Confirming Equilibrium driven by wrong beliefs.¹ The parallel of this finding in the labor market would be a

¹This paper follows Davis and Holt (1993) and Roth (1994) using the following terms:
experiment: the collection of all data;
session: the collection of data involving the same group of subjects on the same day;
treatment: a unique configuration of parameters, variables and rules;

situation in which minority workers, after having been discriminated against, expect that unfavorable conditions continue while biased employers have actually disappeared. Hence, minority workers reduce their effort, and accordingly they are promoted less frequently, but eventually they discover that biased employers have disappeared and balanced promotions across populations of workers are observed.

The strategy method, on the other hand, supports the Self-Confirming Equilibrium driven by wrong beliefs.² In fact, advantaged and disadvantaged subjects react in a different way when *ad hoc* aggregate statistics are displayed. Subjects are asked to bid five times, after five different fictitious distributions of prize winners across populations in the previous period have been displayed. This fictitious distribution shows a fraction of prizes awarded to the Blues decreasing from 80% to 0% in subsequent trials. Although all subjects are informed that there is no computer that awards the prize according to the color label, what happens is that blue players are influenced by the aggregate statistics showing that a decreasing fraction of them gets the prize, and thus bid less and less. In turn, their lower bids make them less likely to win, leading to unequal outcomes, which are consistent with wrong expectations that they were less likely to get the prize.

The structure of the paper is the following. Section ?? describes the theoretical framework behind the experiment and summarizes the testable implications. Section ?? outlines the contributions to the literature that are related to this experiment. Section ?? displays the design of the experiment as well as its procedure. Section 5 contains the results, and Section 6 draws some conclusions.

2 Theoretical framework

This paper aims to provide experimental evidence concerning some testable implications derived from the model that analyzes the role of workers' expectations in explaining observed unequal outcomes in the labor market (see Filippin (2003)). This section provides a summary of the model, emphasizing its testable implications. After the experiment is presented in section ??, several features are contrasted and compared in more detail with the corresponding parts of the model.

The model is formalized as a two-stage game of incomplete information in which populations of workers and employers are involved. In every constituent game, i.e. in every repetition of the game played by agents randomly drawn from their populations, one employer and two workers, one of whom is a minority worker, are randomly matched. The workers choose among three levels of effort (low, intermediate, high) and the employer promotes one and only one of the two workers after having observed their output, which is a function of observable effort and unobservable taste for work. The promotion is desirable because

trial: a decision unit, one repetition of the game.

²The strategy method, developed by Selten (1967) is a procedure that asks a group of subjects to design their strategies after having repeatedly played a game.

the job assigned to the promoted worker is assumed to be characterized by a lower cost of effort. Promotions also depend on employer’s type, which captures the possible disutility of promoting a minority worker (discriminatory tastes).³ Workers know that there are two types of employer, but they do not know whether the employer they face is discriminatory or not. Also the distribution of types within the population of employers is unknown and workers have beliefs about that.

The importance of workers’ expectations can be appreciated by comparing the equilibrium outcome in terms of promotions that may arise when minority workers overestimate the percentage of employers characterized by tastes for discrimination with a situation in which their beliefs are correct *ceteris paribus*. Assuming, for the sake of simplicity, that employers do not discriminate against minorities either directly or statistically, and that all the other sources of heterogeneity such as the distribution of ability among workers have been neutralized, unequal outcomes may still arise due to minority workers’ wrong expectations. In other words, wrong beliefs about being discriminated against may be self-confirming. In this circumstance what happens is that in equilibrium minority groups, who expect being discriminated against, exert less effort on average, because of a lower expected return. This induces a lower percentage of promotions within minority workers, which in turn is consistent with their beliefs that employers are characterized by discriminatory tastes. On the other hand, when beliefs are correct symmetric outcomes are observed.

It is worth stressing that a necessary condition for such a Self-Confirming Equilibrium is that beliefs of majority and minority workers differ. If both groups have wrong but similar beliefs about the fraction of discriminatory employers, their behavior will also be similar, and balanced outcomes in terms of promotions should be expected. The dataset used by Filippin and Ichino (2003) provides interesting evidence that, although men and women share very similar expectations about the magnitude of the gender wage gap, the importance they assign to the underlying causes differs. In fact, while a larger fraction of men think that “actual differences between men and women” matter, a larger fraction of women points towards the “employers’ discriminatory tastes” as one of the causes for the expected gap.

Several implications arising from this model can be tested in the laboratory:

1. *When beliefs are correct, workers’ behavior should not significantly depart from the Bayes-Nash Equilibrium of the game.* In particular:

- when it is common knowledge that there is no discrimination, i.e. when the game is like a symmetric tournament, all workers should exert an inefficiently high level of effort;
- when a known amount of discrimination affects workers’ behavior, there should not be systematic differences across populations. In other words, the

³Observable effort and incomplete information are the main features that distinguish this approach from the tournament literature started by Lazear and Rosen (1981). The two approaches share most of their predictions, one of which being that discrimination, when it is common knowledge, affects the two populations of workers in the same way (see also Section ??).

effort exerted by majority and minority workers should decrease in a similar way.

2. *Workers' behavior should not differ across populations even when beliefs are wrong, provided that the two populations of workers share similar beliefs.*

3. *Workers who overestimate discrimination exert a lower effort than workers characterized by correct expectations.* This is the key mechanism that might drive the labor market towards unequal outcomes even when discriminatory tastes have disappeared. In Filippin (2003) a static framework is used and it is assumed that minority workers are those who might have wrong beliefs. Behind the static model there is an implicit dynamic: minority workers who have experienced direct discrimination for a long period continue to expect of being discriminated against even though discriminatory tastes have disappeared (hysteresis).

4. In the model it is assumed that players have costless access to aggregate outcomes that can be used to form their beliefs. An experiment can test *to what extent subjects' behavior is affected by aggregate information*, as the literature on information cascades suggests.

3 Literature Review

Although the role of workers' expectations in explaining unequal outcomes has never been the focus of experiments, several contributions to the literature are relevant as far as this experiment is concerned. They can be divided into three groups:

1. Discrimination and asymmetric tournaments.
2. Information cascades, sunspot and hysteresis.
3. All-pay auctions.

3.1 Experimental Studies of Discrimination and Asymmetric Tournaments

Experiments closely related to the experiment presented in this paper are those concerning either statistical discrimination or asymmetric tournaments. This subsection concentrates on experiments based on economic factors; a survey of many other experiments based on group identification or status can be found in Anderson, Fryer and Holt (2002).

The literature concerning experimental studies of discrimination is thin and in general not directly related to the experiment presented in this paper, with a few exceptions. Fryer, Goeree and Holt (2002) describe the results of experiments that may produce (and sometimes do) a pattern of experience-based discrimination consistent with the statistical discrimination models proposed by Arrow (1973) and Phelps (1972). Employers have to decide whether to hire or not workers from two otherwise identical populations, "green" and "purple". The hiring decision is affected by an observable test score, which in turn depends on a worker's (unobserved) investment decision, like education or training. The

cost of investing is random and it is set to be systematically higher for the workers of one population during the first ten out of sixty rounds, while from the eleventh onward it drawn from the same distribution. Moreover, players have access to aggregate information, given that the average investment and hiring percentages for the workers of each color are displayed at the end of each round. The authors find that a different average investment emerges, and then a lasting and self-reinforcing mechanism operates in such a way that fewer workers of that group are hired, the fraction of that group of workers investing decreases even further and so on, leading to multiple equilibria with discrimination. There are certain dimensions in which this experiment should be explicitly compared with that presented below. In particular, it is worth noting that, similarly to the experiment described in this paper:

a) there is a real effect (the different distributions from which investment costs are drawn) that is withdrawn during the experiment, but that have long-lasting effects (hysteresis);

b) aggregate statistics are used to convey information to players, although there is no specific treatment focusing on the presence vs. absence of aggregate information.

c) there is an endogenous decision (investing or not) that makes *ex ante* equal populations potentially different in equilibrium.

Hargreaves-Heap and Varoufakis (2002) use the Hawk-Dove game to show that starting from two populations that differ because of a payoff-irrelevant observable characteristic only (red and blue label), different roles associated with different payoffs (i.e. discriminatory conventions) may emerge. Löhmer (2000) finds in a Battle of the Sexes experiment that females are more likely to be discriminated against by other females.⁴

The literature concerning experimental studies of asymmetric tournaments is more established and some papers can fruitfully be used as a benchmark, in particular Schotter and Weigelt (1992) and Bull, Schotter and Weigelt (1987). The former presents an experiment aiming to test the theoretical predictions of the asymmetric tournament theory as presented by O’Keeffe, Viscusi, and Zeckhauser (1984). In particular, they focus on the predicted trade-off between equity and efficiency associated with affirmative actions, finding contrary evidence.

From the theoretical point of view, asymmetric tournaments have many things in common with the model in Filippin (2003). In line with the old saying that different opinions are necessary for a horse race to take place, both involve uncertainty. What distinguishes them is the fact that effort is not perfectly observable in the asymmetric tournament literature, while incomplete information

⁴Other experiments concerning statistical discrimination have been proposed by Davis (1987) and Anderson and Hauptert (1999), both relying on exogenous differences that characterize the two populations. The former finds weak evidence that the larger population has better outcomes. The latter provides evidence that workers belonging to a population characterized by a lower average innate productivity are less likely to be hired, with the likelihood depending on the cost of discovering the individual type. Strictly speaking, it can be argued that the framework of these experiments cannot be classified as discrimination.

about the opponents' type-strategy set characterizes Filippin (2003). Furthermore, the two approaches share most of their predictions, in particular that the behavior of advantaged and disadvantaged workers should change in a similar way when discrimination is common knowledge.

The experiment presented by Schotter and Weigelt (1992) can be used as a benchmark also from the methodological point of view. Three points are particularly relevant. First, the authors want to avoid carryover effects from one treatment to another. Consequently, each subject is allowed to participate in one treatment only. In the present paper, instead, the main goal is to test the existence of hysteresis, and therefore carryover effects are part of the picture. Hence, the treatments are designed in such a way that every subject faces both a symmetric game and situations characterized by discrimination. Second, in the experiment proposed by Schotter and Weigelt (1992) players are matched once and for all within every session. This is more likely to lead to cooperation, or at least to strategic interaction, that would instead disappear with a random matching repeated before every period (see Duffy and Ochs (2003)).

The experiment just described closely follows an earlier experiment by Bull, Schotter and Weigelt (1987), where asymmetries and affirmative actions were not the main focus. It is worth noting that both experiments report a tendency of disadvantaged workers to over-supply effort in uneven tournaments, as if asymmetries elicit greater effort.⁵ Moreover, Schotter and Weigelt (1992) show that advantaged workers make more effort than predicted in unfair tournaments. Section 5 will show evidence that disadvantaged players sometimes over-bid in unfair contests, while some evidence emerges that advantaged players bid more than predicted when the rules of the tournament are particularly favorable.

3.2 All-pay auctions

All-pay auctions are characterized by the fact that bids are given back neither to the winner nor to the loser. The model behind the experiment is related to an all-pay auction, insofar as there is no compensation for the loss of utility that a non-promoted worker suffers when he have exerted an effort higher than the level that would be optimal if promotion was not an issue. In an all-pay auction the prize goes to the highest bidder, so that each player has the incentive to overbid the others, as long as this ensures a positive payoff.⁶ When the value of the prize exceeds the sum of the endowments an equilibrium in pure strategies exists, implying full dissipation of the endowments. Otherwise, in symmetric all-pay auctions, the result that the sum of the expected bids equals the value of the prize is supported by mixed strategies equilibria (Baye, Kovenock and de Vries, 1996). Rational agents never over-dissipate the value of the rent if they have the opportunity to bid zero. However, a relaxation of the rationality via the possibility of decision errors is enough to support a theoretical framework where

⁵Asymmetric contests are defined “uneven” when agents are different, and “unfair” when contestants are identical but the rules favor one of them.

⁶The literature on symmetric rank-order tournaments started by Lazear and Rosen (1981) shares some of the features concerning all-pay auctions.

over-dissipation can be observed (Anderson, Goeree and Holt, 1998) consistently with experimental evidence like Davis and Reilly (1998).⁷

The promotion game without discrimination in Filippin (2003), which is also tested in this experiment, is equivalent to an all-pay auction. Not surprisingly, the Nash equilibrium is a corner solution where endowments are fully dissipated. It might be feasible to increase the value of the endowment (or to decrease the value of the prize) to test also for over-dissipation. However, this goes beyond the goal of this paper.

3.3 Information Cascades, Sunspots and Hysteresis

This experiment is also related to the strand of literature on information cascades (or herd behavior) started by Banerjee (1992) and Bickchandani, Hirshleifer and Welch (1992). Information cascades occur when the initial decisions of other players coincide in a way that it is optimal for each of the subsequent individuals to ignore his or her private signals and follow the established pattern. Particularly interesting as far as this paper is concerned is the sub-case of a reverse cascade (also called bad herd, or lemming type behavior), which happens when the first decision-makers choose the incorrect state of the world, and the followers join the resulting pattern of mistakes despite the fact that their private signals are more likely to indicate the correct state. In Filippin (2003) reverse cascades can justify the Self-Confirming Equilibrium with unequal outcomes when discriminatory tastes have disappeared, since workers access aggregate information that might affect subjects' behavior over and above individual outcomes. Anderson and Holt (1997) provide evidence from the lab of the existence of such information cascades, and, to a minor extent, of reverse cascades. Hey and Morone (2002) show that lemmings survive even within a market contest.

The theoretical sunspot model postulates that agents believe that a variable, which is in fact unrelated to the economy, has real effects, and shows that such beliefs can induce the agents to behave in a manner that provides support for the postulated beliefs. Sunspots were introduced to the laboratory by Woodford (1990), who shows that cyclic sunspot equilibria can asymptotically emerge in an OLG framework when agents follow some adaptive learning schemes. Marimon, Spear and Sunder (1993) do not find evidence that sunspot equilibria exist when the extrinsic variable is not correlated with some real shock. However, they do find evidence that sunspots matter, taking the form of common past experience that influences agents behavior even when the real shock (correlated

⁷When within such a framework subsequent bids are allowed before the prize is assigned, it is easy to observe that bidding spirals out of control, as in the Dollar Auction Game presented by Shubik (1971), where a dollar is awarded to the highest bid. Since the expenditures are sunk, it would be rational to increase a bet whenever doing so increases the expected return more than the amount of the additional bet. There is no stable equilibrium (at least in pure strategies) as long as the endowment of each player exceeds the value of the prize. When one bid exceeds the value of the prize, the motivation of the remaining bidders changes from a desire to maximize returns to one of minimizing losses. Thus, the question transforms from "How much can I win?" to "How do I keep from losing?" and escalation is easily observed, like in the classroom experiments described by Murnighan (2001).

with sunspots) has been removed. This is a combination of hysteresis, i.e. the lagging of an effect behind its cause, and sunspots. What the experiment in section ?? tries to figure out is the existence of hysteresis without sunspots. In this case discriminatory tastes are the key variable that has real effects and that is withdrawn, while there is no extrinsic signal that drives the behavior of agents after the real shock disappears.

The paper by Bull, Schotter and Weigelt (1987), already mentioned when talking about asymmetric tournaments, is relevant also as far as the role of information is concerned. Giving contestants additional, but not complete, information about the actions of their opponents appears to slow the rate at which agents converge on their optimal choice. Section 5 will show that no significant differences are associated with the behavior of players accessing aggregate statistics in addition to individual outcomes.

4 The Experiment

The game captures the main features of the model that explores the role of workers' expectations in explaining observed unequal outcomes, presented in Filippin (2003). The game is much simpler than the model in order to be easily played. At the same time, the subjects are not made aware of the underlying economic relations being tested. Thus, keywords like discrimination, labor market, employer, worker, male and female are never used. This minimizes the risk that idiosyncrasies might enter the experiment and confound the results.

4.1 Sketch of the game

Participants are randomly divided into two populations: red and blue. In every trial every participant has an endowment of 10 Euro cents and can decide how many cents to allocate as a sort of lottery ticket to get a prize worth 25 Euro cents. Bets are not given back to the players, neither to the winner nor to the loser, making the game equivalent to an all-pay auction. Therefore, at the end of the trial the winner gets 25 cents plus the amount not bet, while the loser only the amount not bet.

In every lottery there are only two participants, one from the red population and one from the blue population. The players know that they face only one opponent within every trial. Subjects are warned that it is possible to face the same opponent more than once during every session, but of course they do not know when, given that random assignment takes place at the beginning of each trial. The prize is awarded to the higher bid and it is split if bids are equal, unless the opponents are assigned to a "crazy computer," which instead assigns the prize to the red player regardless of the bids. The fraction of crazy computers, and whether workers know it or not, vary across treatments (see section 4.3).

4.2 Contrast and comparison with the model

Similarities with the model in Filippin (2003) are straightforward only if one knows that this model is what the experiment aims to test. Colors (red and blue) are the equivalent of gender (or race, etc.). The endowment of 10 cents is the same as the utility level when intermediate effort, i.e. the optimal level of effort when promotions are not an issue, is exerted. The amount bet plays the role of additional effort exerted to enhance the probability of being promoted. The prize stands for the promotion and, finally, the “crazy computers” play the role of discriminatory employers. The game is played under different parameter settings and information structures (see section 4.3).

4.2.1 Populations and Number of Types

As already mentioned, red and blue labels are the equivalent of the payoff-irrelevant observable characteristic that distinguishes minority from majority workers. The color label is assigned randomly to every participant and lasts for the whole experiment.

The role of the population of employers is played by the computers, which implements the employers’ equilibrium strategies in the model. The crazy computers never assign the prize to the members of the blue population. The “fair computers” instead assign the prize to the player who made the higher offer and they split the prize when bids are equal. Hence, only the blue players risk being discriminated against.

In the theoretical model it is necessary to assume that workers are of different types for the employers to have some uncertainty about their productivity in the second period. In the experiment the distinction of different types would make the game much more complicated, given that subjects are not familiar with the concept of payoff-type. A further appreciable gain in simplicity is that, since the computers directly play the equilibrium strategy of the employers, it is unnecessary to play the second stage of the theoretical model.

4.2.2 Utility function and Nash Equilibria

The utility function used in the model is not implemented directly in the experiment because it would be quite cumbersome to deal with it in the limited time-spell of the experiment (about 75 minutes). However, the game sketched in section ?? implies a simplified but very close version of it. In both cases players have the opportunity to give up some utility with certainty in exchange for an uncertain but higher return. In the model, supplying a high effort is a sub-optimal decision considering the first period only, but the loss of utility can be more than counterbalanced if the worker is promoted, since the job assigned to the promoted workers is characterized by a lower cost of effort. On the other hand, if the worker is not promoted (s)he suffers a net loss of utility with respect to if where (s)he had chosen the “safe” option, i.e. intermediate effort. The risk of not being promoted has its counterpart in the possibility to

Table 1:
Payoffs for the two players when no computer prefers Reds (Euro cents)

red;blue	0	1	2	...	9	10
0	22.5; 22.5	10; 34	10; 33	...	10; 26	10; 25
1	34; 10	21.5; 21.5	9; 33	...	9; 26	9; 25
2	33; 10	33; 9	20.5; 20.5	...	8; 26	8; 25
...
9	26; 10	26; 9	26; 8	...	13.5; 13.5	1; 25
10	25; 10	25; 9	25; 8	...	25; 1	12.5; 12.5

bid without getting the prize in the experiment. However, also in this case the player has a “safe” option, which is bidding zero.

Table 1 summarizes the payoffs for the two players given the decision rule of a fair computer who assigns the prize to the subject bidding more and splits the prize when the bids are equal. For instance, the cell 22.5; 22.5 is associated with both subjects bidding zero and the prize being split. The only Nash equilibrium occurs when both players bid 10 cents. Hence, full dissipation of the endowments is predicted, not surprisingly given that this game is equivalent to an all-pay auction with the value of the prize exceeding the sum of the endowments.

This Nash equilibrium is not efficient, because both players would be better off playing 0, but in that case both would have the incentive to deviate bidding a positive quantity. This is similar to what happens in the theoretical model for the combination of parameters that displays a Bayes-Nash Equilibrium with both workers supplying a high effort. Given that the model imposes that one and only one worker is promoted, if both workers offered an intermediate effort, the probability of being promoted would be the same and both would have a higher utility in the first period. However, neither worker would be maximizing his/her utility because (s)he has a profitable deviation supplying a high effort.

Table 2 summarizes the payoffs for the two players given the decision rule of a crazy computer who always assigns the prize to the red player regardless of players’ bids. Playing 0,0 is a dominant strategy for both, just as intermediate effort is a dominant strategy in the model when there are discriminatory employers only.

When there are both crazy and fair computers, payoffs are a linear combination of these two matrices, using as weight the fraction of crazy computers. When this fraction is unknown the equivalent of the expectation-driven Self-Confirming Equilibrium characterized by unequal outcomes may arise, as long as crazy computers have disappeared but the minority group still believes there are some.

An important difference with respect to the model is that in the experiment bids are roughly continuous, while in the model the set of choices is restricted to low, intermediate and high effort. A roughly continuous set of choices is likely to reduce the probability of observing the Self-Confirming Equilibrium driven by wrong beliefs, since the cost of experimenting to discover the true state of

Table 2:
 Players' payoffs when facing a PC that prefers Reds (Euro cents)

red;blue	0	1	2	...	9	10
0	35; 10	35; 9	35; 8	...	35; 1	35; 0
1	34; 10	34; 9	34; 8	...	34; 1	34; 0
2	33; 10	33; 9	33; 8	...	33; 1	33; 0
...
9	26; 10	26; 9	26; 8	...	26; 1	26; 0
10	25; 10	25; 9	25; 8	...	25; 1	25; 0

nature becomes very low (see also section 5).

4.3 Design and Procedure of the Experiment

One aspect of the experimental procedure needs to be stressed. All the treatments were proposed within each of the seven sessions of the experiment. Hence, all the subjects played facing the whole set of parameters. This procedure implies potential carryover effects from one parameter set to the others, as well as confounding factors arising because of framing, learning and fatigue. However, testing the existence of carryover effects (hysteresis) is one of the primary goals of the paper, and therefore such an approach is necessary. Moreover, an econometric approach to the analysis of the data allows us to control for any observable and/or unobservable individual characteristic that might affect the choices of the participants during the experiment, including framing, learning and fatigue. To minimize the role of confounding factors, simultaneous parameter changes are avoided.

One of the testable implications concerned the role of information. For this reason in two sessions (3 and 4) subjects get information on all the outcomes within the session, in addition to their own outcomes, when the fraction of crazy computers is unknown (Treatments 2 and 6).

What follows is a sketch of the rules and the procedure of the experiment, which has been run using the zTree software.⁸ We recruited subjects from undergraduate courses at the University of Milan. Most of the subjects were inexperienced. Participants were first randomly assigned numbers and seats. Subjects were told that their physical identity was not associated with their choices during the experiment, the subjects' numbers being their personal identification. They were given written instructions that were also read aloud by the experimenters, stressing that the amount they earned was a function of their decisions. In addition, instructions were also displayed on the screen at the beginning of each treatment.

Quiz 1. After questions are raised by subjects, a quiz is run to test their comprehension of the game. Given three different pairs of bids, they are asked

⁸The zTree software was developed at the University of Zurich, Institute for the Empirical Research in Economics (see Fischbacher, 2002).

to answer who is awarded the prize and to compute their earnings. If wrong answers are given, the subsequent screen shows the subject the correct answer and, in the case of earnings, the way to compute them. Subjects are invited again to ask questions about anything that is not clear.

Assignment to red or blue population. The color of the population is then randomly assigned to every participant by means of an algorithm, in such a way that unobserved and uncontrolled characteristics are not correlated with the focus variables. The color is assigned once and for the whole experiment.

Treatment 1 - 10 trials - Random matching of each Blue with a Red, with the possibility of facing the same opponent more than once. There is no “crazy computer” and players know it, i.e. the game is equivalent to an all-pay auction. Each subject is asked to bid from 0 to 10 Euro cents in order to get the prize. The goal of this treatment is twofold. On the one hand, it makes subjects familiar with the game. On the other hand, it tests whether players cluster around the unique Nash Equilibrium with both players bidding 10 cents, i.e. fully dissipating their endowment. Payoffs are reported in Table 1. After participants have decided the computer displays:

- a) how much the two opponents bid;
- b) who wins the prize;
- c) individual earning.

Introduction of crazy computers. After Treatment 1, the participants are told that some crazy computers, i.e. computers that assign the prize to the member of the red population regardless of the amounts bid, will be introduced into the game. During each treatment they might face either a fair or a crazy computers, but obviously they are not told at the beginning of each trial whether the computer running that trial is crazy or not. It is made clear that in every trial there is a random matching with an opponent of the other population as well as with a PC. Hence, the probability of facing a crazy computer is the same in every trial of each treatment and does not depend on the type of computer faced in previous trials.

Participants are warned that the maximum attention has been paid in order that every subject has the same expected reward. In particular, members of the blue population will receive an additional lump sum reward of 3.5 Euro to compensate them for their lower chances of being awarded the prize. Hence, at the end of the experiment everybody will have had the same chance to earn the same expected reward with differences depending on participants’ actions only.⁹

Quiz 2. Subjects are then asked to answer another short quiz to test their comprehension of the game when crazy computers are introduced. Participants are asked:

- a) to answer whether they are assigned the prize or not, and to compute

⁹The lump sum compensation has been introduced to prevent members of the blue population from feeling tempted to hinder the experiment. The amount of the compensation (3.5 Euro) was derived from the difference between the expected average earnings of Reds and Blues, had they followed the theoretical predictions. It turned out that the compensation was higher than the actual difference of earnings arisen during the experiment between Reds (11.05 Euro) and Blues (9.08 Euro).

their net position under four different circumstances in terms of bids or type of computer;

b) to compute their earnings and to infer the type of computer given four different pairs of bids and the associated decision about the prize.

In both cases, if wrong answers are given, the subsequent screen shows the subject the correct answer and, in the case of earnings, the way to compute them. Subjects are invited again to ask questions about anything that is not clear.

Treatment 2 - 15 trials - Random matching of each Blue with a Red, with the possibility of facing the same opponent more than once. There is no crazy computer but in this case players do not know this. They are just told that the fraction of crazy computers can range between 0% and 100% and that it is held constant during the treatment. Each subject is asked to bid in order to get the prize. In this case theoretical predictions are uncertain, since the best replies depend on beliefs and rationalizability does not allow us to delete any strategy. After each participant has decided the computer displays:

- a) how much the two opponents bid;
- b) who wins the prize;
- c) individual earning.

In sessions 3 and 4 subjects access aggregate information as well. They are shown the distribution of bids among red and blue players, together with the percentage of Blues and Reds to whom the prize is assigned. After statistics have been displayed, subjects are finally asked to report their beliefs about the actual fraction of crazy computers. Expectations are elicited implementing a lottery in which each subject has a probability of winning that is correlated with the number of times in which her beliefs are approximately (error <5%) correct.

This treatment is repeated after players experience discrimination.

Treatment 3 - 5 trials - the same as Treatment 1 except that now there is a 10% probability of facing a crazy computer and players know it. Particular emphasis is devoted to explaining that computers are randomly assigned at every trial to make it clear that a subject does not face the same PC during the whole sequence of trials. The goal of this treatment is twofold. On the one hand, it makes subjects familiar with the game when they know that crazy computers are introduced. On the other hand, it tests whether players over-react to the introduction of crazy computers, given that the unique Nash Equilibrium of the game predicts that both players should still bid 10 cents, i.e. fully dissipating their endowment (expected payoffs are reported in Table 3). After each participant has decided, the computer displays the same information to all the players in all sessions, i.e. in this case even in sessions 3 and 4 no one has access to aggregate information but only to individual statistics, and in particular:

- a) how much the opponents offered;
- b) who wins the prize;
- c) individual earning.

Treatment 4 - 5 trials - like Treatment 3 except that now there is a 50% probability of facing a crazy computer and players know it. As Table 4 shows,

Table 3:
Players' payoffs with 10 percent crazy computers (Euro cents)

red;blue	0	1	2	...	9	10
0	23.75; 21.25	12.5; 31.5	12.5; 30.5	...	12.5; 23.5	12.25; 22.5
1	34; 10	22.75; 20.25	11.5; 30.5	...	11.5; 23.5	11.25; 22.5
2	33; 10	33; 9	21.75; 19.25	...	10.5; 23.5	10.25; 22.5
...
9	26; 10	26; 9	26; 8	...	14.75; 12.25	3.5; 22.5
10	25; 10	25; 9	25; 8	...	25; 1	13.75; 11.25

Table 4:
Players' payoffs with 50 percent crazy computers (Euro cents)

red;blue	0	1	2	...	9	10
0	28.75; 16.25	22.5; 21.5	22.5; 20.5	...	22.5; 13.5	22.25; 12.5
1	34; 10	27.75; 15.25	21.5; 20.5	...	21.5; 13.5	21.25; 12.5
2	33; 10	33; 9	26.75; 14.25	...	20.5; 13.5	20.25; 12.5
...
9	26; 10	26; 9	26; 8	...	19.75; 7.25	13.5; 12.5
10	25; 10	25; 9	25; 8	...	25; 1	18.75; 6.25

in this case there is no Nash Equilibrium in pure strategies, because a cycling pattern would emerge. Starting from the Nash Equilibrium of the last trial of the previous treatment (10;10), if red players stick to their previous bid blue players no longer find it convenient to make a positive bid, and therefore they prefer to drop out by bidding zero. This cannot be an equilibrium because red players could get the full prize by bidding 1 instead of 10, but at that point blue players would have a profitable deviation overbidding by one cent the opponent. The same would apply to red players and so on, until the mechanism started again. There is only one prediction: the average bid of Reds and Blues should not significantly differ. After participants decide, the computer displays the same individual statistics as in the previous treatment.

Treatment 5 - 5 trials - equal to Treatment 4 except that now there is a 90% probability of facing a crazy computer and players know it. As Table 5 shows, although there is no Nash Equilibrium in pure strategies, it is possible to predict that both players should bid from 0 to 2 cents, because bidding more than 2 cents is never a best reply to any beliefs. After the participants have decided, the computer displays the usual individual statistics.

Treatment 6 - 15 trials - Repetition of Treatment 2: random matching of each Blue with a Red, with the possibility of facing the same opponent more than once. There is no crazy computer but players do not know it; they are just told that the fraction of crazy computers can range between 0% and 100% and that

Table 5:
 Players' payoffs with 90 percent crazy computers (Euro cents)

red;blue	0	1	2	...	9	10
0	33.75; 11.25	32.5; 11.5	32.5; 10.5	...	32.5; 3.5	32.25; 2.5
1	34; 10	32.75; 20.25	31.5; 10.5	...	31.5; 3.5	31.25; 2.5
2	33; 10	33; 9	31.75; 9.25	...	30.5; 3.5	30.25; 2.5
...
9	26; 10	26; 9	26; 8	...	24.75; 2.25	23.5; 2.5
10	25; 10	25; 9	25; 8	...	25; 1	23.75; 1.25

it is held constant during the treatment. As in Treatment 2, individual statistics are displayed after the participants have decided, while in the same two sessions (3 and 4) subjects access aggregate information as well. This is identified as a different treatment because of carryover effects that might affect subjects' behavior. Average bids that significantly differ from Treatment 2 would signal persistent effects of discrimination. Blues bidding less than Reds would be evidence in favor of the Self-Confirming Equilibrium in the model.

Since the set of choices in the experiment is roughly continuous, the Self-Confirming Equilibrium that may arise differs from the Self-Confirming Equilibrium in the model. As in the model, minority workers should choose the "safe" option, i.e. bidding zero (equivalent of intermediate effort), as long as they think there is a sufficiently high fraction of crazy computers. However, while in the model the majority workers would exert high effort, in the game it is enough to bid 1 cent instead of 10 (which would be the equivalent of high effort) if Reds believe that the Blues play zero. In more detail, the Self-Confirming Equilibrium driven by wrong minority workers' beliefs corresponds in the game to a situation in which:

- Reds bid 1 cent, which is a best reply to beliefs that crazy computers are up to 8% and that Blues bid zero.
- Blues think that at least 92% of the computers are crazy and that Reds bid 1 cent. Given this beliefs bidding zero is a best reply.
- The prize is awarded to the Reds.¹⁰

¹⁰This Self-Confirming Equilibrium is not unique. Another equilibrium is obtained simply flipping the color labels in the example above. However, only the equilibrium in which the prize is assigned to the Reds is consistent with the model. If the equilibrium in which the prize is assigned to the Blues occurs, it would be strong evidence against the model.

On the other hand, if all players believe that at least 92% of the computers are crazy and that the opponent play zero, both Reds and Blues should bid zero. However, in this case a few trials are enough to falsify players beliefs that there are many crazy computers, given that the prize would be split every time.

Finally, another Self-Confirming Equilibrium is supported by both subjects bidding the whole endowment and thinking that there are at most 20% of crazy computers.

Treatment 7 - 5 trials - Repetition of Treatment 3: 10% probability of facing a crazy computer and players know it. Again, this is identified as a separate treatment because of carryover effects that might affect subjects' behavior. Significantly different bids than in Treatment 3 would point toward persistent effects of discrimination. Blues bidding less than Reds would be evidence in favor of the Self-Confirming Equilibrium driven by wrong beliefs.

Strategy method. - 5 trials - There is no crazy computer and players know it. Each subject is asked to bid five times, every time after a different fictitious distribution of prizes between populations in the previous round has been displayed. This fictitious distribution shows a fraction of prizes won by the Blues decreasing from 80% to 0% in subsequent periods. In this case each subject wins the prize if her bid exceeds the average bid of the population of opponents. The individual results of each trial are shown together after the fifth trial, so that they do not affect the choices of the subjects during the treatment. Since there is no crazy computer, and subjects know it, subjects of both populations should dissipate entirely their endowment and pay no attention to the aggregate statistics.

At the end of the experiment a questionnaire is proposed, reminding participants that their physical identity was not associated with their choices and their answers during the experiment. Questions concerned academic as well as personal information. In section 4.4 some descriptive statistics of the pool of subjects are summarized.

70 subjects participated in the experiment, which consisted of seven sessions. The sessions lasted approximately 75 minutes and were composed of a minimum of 8 and a maximum of 12 subjects. Euro cents were the currency used during the experiment. Earnings ranged between 9 and 15.5 Euros (11.8 on average). Average earnings of the Reds were lower than average earnings of the Blues once the lump sum compensation of 3.5 Euros was taken into account. (11.05 Euro vs. 9.08+3.5 Euro, respectively).

4.4 Sample description

From the information collected by means of the final questionnaire, it turns out that males are over-represented in our sample (67% vs. 33%), and that the average age of the pool is about 21 years. Most of the participants (89%) comes from the School of Political Sciences, and is enrolled in the third year of the degree program. The final mark at the exit of secondary school was chosen as a proxy for a student's ability; the variable has been rescaled in the range [0,1]. Two thirds of the sample come from high schools (*licei*) and one fourth from technical schools (*istituti tecnici*). Two specific questions were asked concerning political and religious orientation. An ordered scale from 0 to 5 has been used to ask subjects about their political orientation (0=left; 5=right), without any label on each possible choice. Two thirds of the subjects report themselves as being center-left, i.e. they chose a value from 0 to 2, and 33% center-right. The average choice is 1.97 while the median choice is 2. With respect to religion, the subjects have been asked to choose from three alternatives: "believer and

Table 6:
Treatment 1: average bid (st dev) and distribution of prize winners

Session (Subjects)	Average bid		Kolmogorov- Smirnov test	Prize winners		
	red	blue		red	split	blue
1 (10)	9.82(0.66)	9.70(0.81)	$p:1.000$	6	41	3
2 (10)	9.58(1.39)	9.76(0.72)	$p:1.000$	5	37	8
3 (10)	9.02(2.80)	9.56(1.51)	$p:0.998$	8	37	5
4 (8)	8.95(1.28)	8.57(2.51)	$p:0.579$	11	14	15
5 (10)	7.44(3.92)	8.92(2.55)	$p:0.112$	9	23	18
6 (10)	9.80(1.28)	9.76(0.85)	$p:1.000$	5	43	2
7 (12)	8.62(2.73)	8.63(2.76)	$p:0.928$	14	28	18
<i>pooled</i> (70)	9.02(2.42)	9.27(1.93)	$p:0.942$	58	223	69

churchgoer,” “believer but not churchgoer,” “non-believer.” The proportion of the last occurrence was around one third.

5 Results

From the quizzes it is possible to infer that subjects have a good albeit imperfect comprehension of the game, given that the average number of wrong answers is about 1.4 out of 6 questions in the first quiz, and 2.5 out of 16 in the quiz with crazy computers.

The results confirm the prediction implied by the Nash equilibrium of the game where subjects know that there is no crazy computer (**Treatment 1**), i.e. that they should fully dissipate their endowment and bid 10 cents. In fact, the average bid is 9.14 cents. In the last trial of Treatment 1, 64 out of 70 subjects fully dissipate their endowment. The distribution of prizes is balanced across populations, in line with predictions. The prize is split in 63.7% of the games, while Reds win 16.6% of the games and Blues 19.7%. In the last trial of the Treatment the prize is split in 30 out of 35 games. Table 6 reports the average bid of each population, session by session, pooling all periods. The last row summarizes the same statistics pooling all the sessions. Blues display a propensity to bid more than Reds in session 5, although the difference is not statistically significant at the 10% confidence level. In fact, the Kolmogorov-Smirnov test does not reject the equality of the distribution functions ($p:0.112$). In the other sessions, as well as in the pooled sample, the behavior of red and blue subjects is very similar, in line with the theoretical predictions.

Let’s skip for one moment the analysis of Treatment 2. When a known and small fraction (10%) of crazy computers is introduced (**Treatment 3**), some subjects slightly over-react (see Table 7). In fact, while the small fraction of crazy computers does not affect the prediction of full dissipation of the endowment, the average bid decreases to 8.3 cents with a higher variability. This is particularly evident in session 3 and, to a minor extent, in sessions 2, 4, 6 and 7.

Table 7:
Treatment 3: average bid (st dev) and distribution of prize winners

Session (Subjects)	Average bid		Kolmogorov- Smirnov test	Prize winners		
	red	blue		red	split	blue
1 (10)	9.60(0.66)	9.88(0.60)	$p:1.000$	2	22	1
2 (10)	7.20(4.20)	9.32(2.08)	$p:0.285$	7	11	7
3 (10)	7.16(4.33)	6.40(4.25)	$p:0.710$	11	7	7
4 (8)	8.45(2.43)	7.75(3.93)	$p:0.832$	7	5	8
5 (10)	7.88(3.26)	8.76(2.17)	$p:0.915$	9	8	8
6 (10)	9.60(1.38)	8.04(3.79)	$p:0.915$	5	19	1
7 (12)	8.27(3.49)	7.83(3.58)	$p:0.958$	13	12	5
<i>pooled</i> (70)	8.30(3.29)	8.29(3.28)	$p:1.000$	54	84	37

Table 8:
Treatment 4: average bid (st dev) and distribution of prize winners

Session (Subjects)	Average bid		Kolmogorov- Smirnov test	Prize winners		
	red	blue		red	split	blue
1 (10)	4.80(4.72)	3.72(4.42)	$p:0.475$	17	5	3
2 (10)	5.20(4.06)	4.32(4.12)	$p:0.475$	22	2	1
3 (10)	1.12(1.39)	5.96(3.77)	$p:0.000$	13	0	12
4 (8)	2.65(3.72)	4.40(3.91)	$p:0.034$	15	2	3
5 (10)	4.28(3.32)	4.72(4.01)	$p:0.710$	22	1	2
6 (10)	6.76(4.32)	5.84(4.52)	$p:0.915$	16	6	3
7 (12)	3.80(4.02)	4.77(3.94)	$p:0.071$	18	3	9
<i>pooled</i> (70)	4.12(4.09)	4.83(4.11)	$p:0.313$	123	19	33

Reds and Blues display a pattern that is not significantly different in each session or in the pooled sample, in line with predictions. Unequal outcomes emerge as a consequence of the crazy computers, with Blues winning only 21.1% of the prizes against 30.8% of the Reds.

When the fraction of crazy computers increases to 50% in **Treatment 4**, the average bid decreases sharply to 4.47 cents (4.12 the Reds, 4.83 the Blues), with Blues bidding clearly more than Reds in most of the sessions. The Kolmogorov-Smirnov test rejects the equality of distribution functions in session 3 ($p=0.000$), 4 ($p=0.034$) and 7 ($p=0.071$), while the difference is not statistically significant in the other sessions or in the pooled sample (see Table 8). The effect of the high fraction of crazy computers is overwhelmingly stronger than the effect of the higher offers of blue subjects in determining the distribution of prize winners (70.2% Reds Vs 18.8% Blues).

When in **Treatment 5** 90% of computers are crazy, the average bid (1.52 cents) stays within the predicted range of [0,2] cents, but there is a lot of variability across sessions. As Table 9 shows, red subjects bid on average a quantity

Table 9:
Treatment 5: average bid (st dev) and distribution of prize winners

Session (Subjects)	Average bid		Kolmogorov- Smirnov test	Prize winners		
	red	blue		red	split	blue
1 (10)	0.96(1.65)	0.92(2.16)	$p:0.915$	24	0	1
2 (10)	3.88(3.55)	1.28(2.11)	$p:0.015$	24	0	1
3 (10)	0.36(0.56)	2.00(2.19)	$p:0.002$	23	2	0
4 (8)	0.15(0.67)	1.85(2.87)	$p:0.000$	17	0	3
5 (10)	0.72(0.68)	0.80(1.44)	$p:0.996$	24	1	0
6 (10)	3.08(4.14)	1.44(2.61)	$p:0.710$	23	0	2
7 (12)	2.07(3.68)	1.53(2.10)	$p:0.135$	29	0	1
<i>pooled</i> (70)	1.66(2.94)	1.39(2.23)	$p:0.203$	164	3	8

that cannot be a best reply to any beliefs in sessions 2, 6 and 7. Reds bid significantly more than Blues in session 2, while the opposite happens in sessions 3 and 4. The difference is not statistically significant in the other sessions or in the pooled sample ($p=0.203$). Unequal outcomes become particularly severe, with Reds winning the prize 93.7% of the times, against 4.6% of the Blues.

Summarizing, subjects follow rather well the theoretical predictions when the fraction of crazy computers is known. This is particularly true in session 1, where deviations from the predicted behavior are negligible. Some departures from the predicted behavior are worth noting, however. In sessions 3 and 4 Blues have the propensity to bid significantly more than Reds when a sufficiently large fraction of crazy computers is introduced. This happens also in one case in session 7. In session 2 red players react to the presence of a negligible fraction of crazy computers. Also in session 6 (blue players) and in session 3 (all players) there is evidence of over-reaction when the fraction of crazy computers is negligible. Reds offer on average an amount higher than rationalizable in sessions 2, 6 and 7 when many crazy computers are introduced.

The analysis of **Treatment 2**, the first with the fraction of crazy computers equal to zero but unknown to the players, shows a high variability both between and within sessions. This is not surprising, given that any bid can be a best reply given some beliefs about the fraction of crazy computers and the opponent's strategy. What turns out to be interesting is the fact that disadvantaged players bid on average significantly more than the opponents in three out of seven sessions (3, 5 and 7). In session 7 this happens because only red players significantly react to the possibility that crazy computers are introduced after Treatment 1, while blue players do not. In session 5 both populations significantly reduce their bids in a similar way, but Blues offer more than Reds also in Treatment 1, when it is known that there was no crazy computer. In session 3 both populations significantly reduce their bids, but Reds more than Blues. In a nutshell, the announcement that there is the possibility of facing crazy computers significantly affects the behavior of most of the subjects, but in some

Table 10:
Treatment 2 - all periods: average bid (st dev) and distribution of prize winners

Session (Subjects)	Average bid		Kolmogorov- Smirnov test	Prize winners		
	red	blue		red	split	blue
1 (10)	8.97(2.48)	8.07(3.90)	$p:0.518$	14	46	15
2 (10)	7.16(4.03)	7.72(3.31)	$p:0.395$	14	39	22
3 (10)	5.90(4.62)	6.81(3.62)	$p:0.010$	26	19	30
4 (8)	5.93(4.25)	5.01(4.22)	$p:0.375$	28	8	24
5 (10)	5.21(4.06)	6.67(3.50)	$p:0.042$	27	7	41
6 (10)	9.07(2.04)	8.69(2.68)	$p:0.996$	18	44	13
7 (12)	6.70(4.36)	7.78(3.56)	$p:0.081$	22	38	30
<i>pooled</i> (70)	7.01(4.04)	7.33(3.68)	$p:0.030$	149	201	175

Table 11:
Treatment 2 - first 5 and last 5 periods: average bid (st dev)

Session (Subjects)	Beginning of Treatment 2		End of Treatment 2	
	Average bid		Average bid	
	red	blue	red	blue
1 (10)	7.88(3.01)	6.60(4.66)	9.96(0.20)	9.60(2.00)
2 (10)	6.92(4.38)	7.36(5.39)	7.72(3.77)	8.04(3.31)
3 (10)	6.28(4.70)	7.28(3.21)	7.00(4.39)	6.32(3.91)
4 (8)	6.15(4.46)	4.35(4.16)	5.60(4.33)	5.20(4.25)
5 (10)	3.08(3.29)	6.04(3.06)	6.56(3.90)	7.60(3.58)
6 (10)	8.00(2.55)	7.68(3.56)	9.80(1.00)	9.32(1.89)
7 (12)	5.73(4.33)	7.10(3.89)	7.47(4.30)	8.13(3.51)
<i>pooled</i> (70)	6.28(4.13)	6.70(3.82)	7.78(3.76)	7.83(3.52)

cases Reds react more than Blues.¹¹

The fraction of crazy computers being unknown, it is interesting to see what happens when the analysis is restricted to the last 5 periods of the treatment, when subjects have the possibility of learning from their experience (and in sessions 3 and 4 also from aggregate statistics) the true state of nature, i.e. that there is, in fact, no crazy computer. Data reported in Table 11 show that learning takes place. In all the sessions the average bid is higher at the end of the treatment with the exception of blue players in session 3 and red players in session 4, the two sessions where aggregate information is available. However, evidence is not strong enough to claim that players in these two cases are learning the “wrong” self-confirming equilibrium with discrimination.

¹¹The Kolmogorov-Smirnov test does not reject at 10% confidence the invariance of players’ behavior between treatments 1 and 2 only in sessions 1 and 6 and only for blue players in session 7. These statistics are not reported in detail in order to save space, but are available from the author.

Table 12:
Treatment 6 - all periods: average bid (st dev) and distribution of prize winners

Session (Subjects)	Average bid		Kolmogorov- Smirnov test	Prize winners		
	red	blue		red	split	blue
1 (10)	9.64(1.70)	9.08(2.77)	$p:0.996$	9	63	3
2 (10)	7.21(4.04)	9.67(1.18)	$p:0.002$	6	43	26
3 (10)	7.08(4.45)	6.12(3.77)	$p:0.003$	31	12	31
4 (8)	5.33(3.99)	5.20(3.95)	$p:0.047$	27	7	26
5 (10)	6.35(4.35)	8.01(3.10)	$p:0.016$	16	25	34
6 (10)	9.96(0.35)	9.63(1.50)	$p:0.996$	5	69	1
7 (12)	9.32(2.39)	8.97(2.70)	$p:0.948$	14	69	7
<i>pooled</i> (70)	7.96(3.69)	8.20(3.23)	$p:0.050$	108	299	118

Table 13:
Treat. 2, last 5 periods; Treat. 6, first 5 and last 5 periods: average bid (st dev)

	End of Treatment 2 Average bid		Beginning of Treat. 6 Average bid		End of Treatment 2 Average bid	
	red	blue	red	blue	red	blue
1	9.96(0.20)	9.60(2.00)	9.32(2.17)	8.00(4.08)	9.60(2.00)	9.40(2.29)
2	7.72(3.77)	8.04(3.31)	8.20(3.49)	9.32(1.84)	6.80(4.40)	10.00(0.00)
3	7.00(4.39)	6.32(3.91)	7.04(4.49)	5.20(3.62)	6.40(4.90)	6.80(3.85)
4	5.60(4.33)	5.20(4.25)	5.50(3.90)	3.75(3.49)	5.30(4.23)	6.30(4.01)
5	6.56(3.90)	7.60(3.58)	5.52(4.38)	6.96(3.51)	6.36(4.59)	8.96(2.32)
6	9.80(1.00)	9.32(1.89)	9.88(0.60)	9.20(2.00)	10.00(0.00)	9.68(1.60)
7	7.47(4.30)	8.13(3.51)	8.63(3.19)	9.00(2.46)	10.00(0.00)	8.53(3.34)
	7.78(3.76)	7.83(3.52)	7.82(3.69)	7.50(3.49)	7.91(3.85)	8.58(3.01)

Let us move to the repetition of the same treatment after subjects have faced an increasing fraction of crazy computers (10%, 50% and 90%), summarized in Table 12. In sessions 2 and 5 Blues offer significantly more than Reds, while in session 3 Reds bid significantly more than Blues.

One of the key questions this paper tries to answer is how subjects' behavior differs before and after having experienced the experiment's version of discriminatory tastes. In sessions 5, 6 and 7 subjects bid significantly more after having faced crazy computers than before. This also happens to Reds in session 1 and to Blues in session 2, respectively. In the other cases significant differences do not emerge, but a significant decrease of bids is never observed. At first glance, this finding is rather puzzling given that the opposite should be expected following the predictions in Filippin (2003). However, as Table 13 shows, restricting our attention to the beginning of **Treatment 6** as opposed to the end of treatment 2, it turns out that in sessions 1, 3 and 4 Blues reduce

their bids, while Reds do not. In theory, this should lead to the Self-Confirming Equilibrium driven by wrong beliefs to emerge, but this effect on disadvantaged subjects' behavior is short-lived and vanishes during the treatment. In fact, at the end of the treatment red and blue subjects behave in a very similar way in most of the sessions.

Hence, evidence about hysteresis is rather weak, because the finding that, after having been discriminated against, Blues bid less than Reds and receive a lower fraction of prizes is not long-lasting. Although a Self-Confirming Equilibrium with wrong beliefs is something that can happen but cannot be expected to emerge regularly, evidence is not fully satisfactory if evaluated from the point of view of the testable implication 3 above. The main reason the experiment does not provide strong evidence in favor of unequal outcomes driven by subjects' expectations, is that the experiment failed to separate the beliefs of the two populations. While in the theoretical model different expectations about the fraction of discriminatory employers are a necessary condition for the Self-Confirming Equilibrium to be observed, in the lab Blues never expect a clearly higher fraction of crazy computers than Reds. A possible explanation for this finding is that the experimental design overemphasized the discontinuity between treatments, preventing carryover effects from entering the picture. Another possible explanation could be that the set of choices was roughly continuous. This induces a very slight difference in the optimal behavior of advantaged and disadvantaged subjects, and therefore also a very low cost of experimenting to discover the true state of nature. Were the choice set dichotomized (e.g. 0, 5 and 10 cents), it would be more costly to experiment and the Self-Confirming Equilibrium would be more likely to emerge.

From **Treatment 7**, i.e. the repetition of treatment 3 after subjects face an increasing fraction of crazy computers, subjects' behavior does not change, with the exception in session 4 of Reds who significantly reduce their bid.¹²

Quite surprisingly, given the evidence from the comparison of Treatments 2 and 6, the strategy method shows evidence supporting the Self-Confirming Equilibrium driven by wrong beliefs. In fact, advantaged and disadvantaged subjects react in a different way when *ad hoc* aggregate statistics are displayed showing a fraction of prizes won by the Blues decreasing from 80% to 0% in subsequent periods. The subjects are told that there is no crazy computer. Hence, subjects of both populations should pay no attention to the aggregate statistics and entirely dissipate their endowment. On the contrary, as the fictitious distribution of prizes becomes less and less favorable to blue players, their offers decrease even though they know that the likelihood of winning the prize does not change, given that there is no crazy computer. On the other hand, red players do not change their behavior significantly (see Table 14).

As a result, the fraction of blue players bidding more than the average of their opponents also decreases from 74.3% in the first period of the strategy method, when the distribution of prizes was supposed to be more favorable, to

¹²These statistics are not reported in detail in order to save space, but are available from the author.

Table 14:

Strategy method: average bid (st. dev) across populations period by period

		Percentage of prizes announced to be won by Blues				
	pop	80	60	40	20	0
1	red	6.80(4.20)	8.00(2.24)	8.80(1.00)	9.00(1.58)	8.60(2.86)
1	blue	8.80(1.98)	8.60(2.00)	4.00(3.82)	3.40(4.39)	0.00(0.00)
2	red	8.80(1.98)	9.00(1.29)	7.60(3.96)	3.20(3.32)	4.20(4.85)
2	blue	9.40(1.22)	9.00(1.29)	6.40(3.07)	2.20(4.00)	2.00(4.85)
3	red	4.20(4.85)	5.20(4.58)	7.20(4.00)	7.20(4.00)	4.40(4.68)
3	blue	7.80(1.75)	8.00(1.71)	7.80(2.45)	6.80(4.00)	3.60(4.55)
4	red	7.00(4.23)	5.75(3.43)	5.00(3.16)	5.75(3.73)	4.25(4.49)
4	blue	8.00(2.05)	8.25(1.97)	4.50(4.26)	4.00(4.35)	3.75(4.25)
5	red	6.20(4.40)	8.80(2.45)	8.60(1.53)	9.40(0.82)	7.60(3.57)
5	blue	7.20(3.73)	10.00(0.00)	8.00(1.94)	6.00(5.00)	4.20(4.85)
6	red	8.00(4.08)	7.60(3.27)	9.80(0.40)	10.00(0.00)	10.00(0.00)
6	blue	8.20(3.67)	7.40(2.55)	7.60(3.96)	2.60(3.85)	4.20(4.85)
7	red	6.00(4.23)	8.17(2.65)	7.33(3.60)	9.50(0.77)	7.00(4.35)
7	blue	7.16(3.44)	7.50(2.33)	4.67(2.86)	3.67(4.57)	3.17(4.00)
<i>pooled</i>	red	6.69(4.25)	7.57(3.22)	7.83(3.15)	7.83(3.33)	6.66(4.37)
<i>pooled</i>	blue	8.06(2.83)	8.37(2.04)	6.14(3.57)	4.09(4.54)	2.97(4.27)

25.7% in the fifth period without any change in fundamentals. What happens in all sessions is that blue players are influenced by the aggregate statistics showing that none of them has won the prize, inducing them to bid less. In turn, their lower bids make them less likely to win, leading to unequal outcomes that are consistent with wrong expectations that they were less likely to get the prize.

Given the design of the experiment, and in particular that every subject is exposed to the whole set of parameter changes, a regression analysis of the data is certainly informative. The limitations imposed by the very low number of independent observations prevent inference from being reliable. However, the interpretation of a regression as a conditional expectation function is not at all affected by the low number of independent observations and sheds more light on the data. In a multivariate framework, with players' bids as a dependent variable, where fixed effects control for any observable or unobservable individual characteristic, also session fixed effects display a high heterogeneity. The fraction of crazy computers, or individual beliefs when the fraction of crazy computers is not known, is obviously the most important variable in explaining the variation of bids, accounting for a 0.61 cents lower bid every 10% of crazy computers. It is not, however, the only one. The dummy variable for the population shows that, everything else being equal, the blue players have a propensity to bid much more (3.5 cents). As far as the learning and framing effects, bids tend to increase within treatments where the fraction of crazy computers is not announced, while across treatments a U-shaped negative effect emerges, with a

minimum in treatment 4 where bids are *ceteris paribus* about 1.3 cents lower than in treatments 1 and 7.

6 Conclusions

This paper is aimed at testing the predictions of a model that explores the role of workers' expectations of being discriminated against as an original explanation for the puzzling long-run persistence of observed discrimination against some minorities in the labor market. The model, presented in Filippin (2003), provides a theoretical framework based on a two-stage game of incomplete information where preferences and beliefs of both sides of the labor market matter. In every constituent game two workers, one of whom is a minority worker, are drawn from their *ex ante* identical populations and randomly matched with one employer.¹³ At the end of the first period the employer promotes one (and only one) worker after having observed the output they have produced, which is a function of their observable effort. Promotions also depend on employer's type, unknown to the workers, which captures the possible disutility of promoting a minority worker. The importance of workers' expectations can be appreciated by comparing the distribution of promotions across populations that arises when minority workers overestimate the percentage of employers characterized by tastes for discrimination with a situation in which such beliefs are correct *ceteris paribus*. This difference becomes crystal clear when there are actually only employers who do not discriminate against the minority either directly or statistically. Even in this circumstance unequal outcomes may emerge, caused by wrong beliefs of being discriminated against that are self-confirming. Minority groups who expect of being discriminated against exert a lower effort on average, because of a lower expected return. This induces a lower observed percentage of promotions within minority workers, which in turn is consistent with their beliefs that there are employers characterized by discriminatory tastes.

The experiment replicates the model using a game where participants are randomly divided into two populations: red and blue. In every trial each participant has an endowment of 10 Euro cents and can decide how many cents to allocate as a sort of lottery ticket to get a prize worth 25 Euro cents. Bets are not given back to the players, neither to the winner nor to the loser, making the game equivalent to an all-pay auction. Therefore, at the end of the trial the winner gets 25 cents plus the amount not bet, while the loser only the amount not bet. In every lottery there are only two participants, one from the red population and one from the blue population. The players know that they face only one opponent in every trial and that it is possible to face the same opponent more than once during the same treatment, but of course they do not know when, given that random assignment takes place at the beginning of each trial. The prize is awarded to the higher bid and it is split if bids are equal, unless

¹³ What distinguishes the population of minority workers is an observable characteristic not related to their productivity (e.g. race, gender).

the opponents are assigned to a crazy computer which instead awards the prize to the red player regardless of the bids.

The mechanism underlying the Self-Confirming Equilibrium driven by wrong beliefs in the theoretical model is tested comparing the outcomes of an identical treatment proposed to the subjects both before and after they face an increasing fraction of crazy computers. A reduction of bids made by disadvantaged subjects is observed in three out of seven sessions, in line with the theoretical predictions, but it vanishes rather quickly during the treatment, failing to generate the Self-Confirming Equilibrium driven by wrong beliefs about discrimination. The parallel of this finding in the labor market would be a situation in which minority workers, after having been discriminated against, expect that unfavorable conditions continue while biased employers have actually disappeared. Hence, minority workers reduce their effort, and accordingly they are promoted less frequently, but eventually they discover that biased employers have disappeared and balanced promotions across populations of workers are observed.

The main reason that the experiment does not provide strong evidence in favor of unequal outcomes driven by subjects' expectations, is that the experiment failed to separate the beliefs of the two populations. While in the theoretical model different expectations about the fraction of discriminatory employers are a necessary condition for unequal outcomes to emerge, in the lab Blues never expected a clearly higher fraction of crazy computers than did Reds. A possible explanation for this finding is that the experimental design laid too much emphasis on the discontinuity between treatments, preventing carryover effects from emerging.

Another possible reason that the experiment does not provide strong evidence in favor of unequal outcomes driven by subjects' expectations, is the fact that the set of choices is roughly continuous. This continuity means that there is only a very slight difference in the optimal behavior of the advantaged and the disadvantaged subjects, and therefore also a very low cost of experimenting to discover the true state of nature. Finally, it is also worth noting that, from a regression where all observable and unobservable individual characteristics as well as the parameters of the experiment are controlled for, Blues display a propensity to bid much more than Reds *ceteris paribus* in spite of the random assignment of the color.

The strategy method, on the other hand, supports the Self-Confirming Equilibrium driven by wrong beliefs. In fact, advantaged and disadvantaged subjects react in a different way when *ad hoc* aggregate statistics are displayed. A fictitious distribution shows a fraction of prizes won by the Blues decreasing from 80% to 0% in subsequent trials. Since all subjects are informed that there is no crazy computer, it might be expected that members of both populations ignore the aggregate statistics displayed and dissipate their endowment entirely. Red players, in fact, do not change their behavior significantly. Blue players, on the other hand, as the fictitious distribution of prizes becomes less and less favorable to them, lower their offers. As a result, the fraction of blue players bidding more than the average of their opponents also decreases from 74.3% in the first trial, when the distribution of prizes was supposed to be more favorable, to 25.7% in

the last trial without any change in fundamentals. What happens in the last trial is that blue players are influenced by the aggregate statistics showing that none of them got the prize, and thus bid less. In turn, their lower bids make them less likely to win, leading to unequal outcomes that are consistent with wrong expectations that they were less likely to get the prize.

The experiment also deals with the relevance of the information structure by dividing the two populations in two subgroups. The first observes individual outcomes only (bids and winners of the games in which the player is directly involved), while the second is also informed about the distribution of bids within, as well as the distribution of promotion across, populations. From the experiment there is no evidence of different patterns between these two subgroups.

Concluding, findings of the experiment provide some evidence supporting the Self-Confirming Equilibrium driven by wrong beliefs, but the evidence cannot be considered fully satisfactory. In order to provide a more robust test of the theoretical model, future experiments should be modified to include such things as a dichotomized choice set or less clear-cut treatments. Another potentially fruitful way to test the model that is left for future research is to try a different version of the experiment in which Reds and Blues may behave differently even though they share similar beliefs about the fraction of crazy computers. Implementing a session in which two Reds and two Blues compete for one prize could be the way to separate directly the behavior of Reds and Blues. In fact, after having observed the presence of crazy computers, they might expect in the following treatment that there are some crazy computers even if this is not the case. In the version of the experiment presented in this paper, similar beliefs lead to similar behavior. If there are two players from each population, instead, even though all players share similar beliefs about the fraction of crazy computers, Reds should still bid a lot because, even though blue players are discriminated, they need to compete against each other in order to get the prize, while for the Blues is pointless to try hard. Once behavior differs, a self-reinforcing mechanism could operate in which Blues bid less, they are not awarded the prize and they bid less and less, while Reds keep bidding significantly more.

Appendix: Instructions

The Experiment - part A (shown at the beginning)

The experiment will last approximately 60 minutes, but the actual length depends on the speed of the slowest participant. The experiment is composed by two quizzes, eight stages and a questionnaire.

Numbers during the experiment represent Euro cents. Your final earnings will be the sum of all the Euro cents you earned throughout the experiment. Earnings depend on your choices as well as on the choices of your opponents during the game that will start in a few minutes.

The game consists of an auction, in which you have to bid in order to get a prize. The game will be repeated several times under potentially different conditions that will be explained at the beginning of each stage.

At the beginning of the experiment an algorithm will assign to every player a color label (red or blue) that will be effective for the whole experiment. The two populations (Red and Blue) will be of equal size. In every repetition of the game your opponent will be an anonymous player randomly drawn from the other population (i.e. if you belongs to the red population you will always play against a Blue and vice versa).

In every repetition of the game you will be endowed with 10 Euro cents. You have to decide how much to bid (from 0 to 10 cents) in order to win a prize worth 25 Euro cents.

- The higher bid wins the prize.
- If bids are equal, the prize is equally split.
- Bids are not given back, neither to the winner nor to the loser.

Your earnings in every repetition of the game depends on two factors:

- 1) The prize: 25 cents if you are awarded the prize, 12.5 cents if the prize is split, 0 if your opponent is awarded the prize;
- 2) How much of your endowment you did not bid.

N.B. You cannot save and transfer money from one repetition of the game to another. If you bid less than 10 cents the amount left will enter your earnings but in the following repetition you will start again with 10 cents.

The Experiment - part B (shown after Treatment 1)

Now a different kind of computers is introduced into the game. These computers, which are called “crazy computers,” award the prize always to the red player regardless of who bids more.

Notice that the computers we are talking about (“normal” vs. “crazy”) are not the computers you have in front of you. The server computer has been programmed to receive the data from the client computers (i.e. to receive the bids that you enter in the PCs in front of you). In every repetition of the game, each pair of players is randomly associated with a partition of the server that can correspond to a normal computer (which assign the prize to the higher bid) or to a crazy computer (which always awards the prize to the Red).

NB: In every repetition of the game, the bids of each pair of players are randomly assigned to a partition of the server computer. Hence, being assigned to a crazy or to a normal computer during one repetition of the game, does not depend on the kind of computer faced in the previous repetition and does not predict anything about the kind of computer that will be faced in the following repetition. From this point of view it is like starting from the beginning at every repetition. You only have to bear in mind that each time there is a percentage of partitions of the server that represent crazy computers.

How many are the crazy computers? The percentage of crazy computers can vary from 0% to 100% during the experiment, but it is held constant within each stage of the experiment. In some stages you will know the fraction of crazy computers, in other stages that fraction will be unknown.

The introduction of crazy computers creates different conditions for Reds and Blues. To ensure that all participants have the same earning opportunity, each Blue will receive a lump sum compensation of 3.5 Euro in addition to what earned during the experiment. This compensation corresponds to the estimated loss induced by crazy computers and DOES NOT depend on the choices you will make during the experiment.

References

- [1] ANDERSON, D.M. AND M.J. HAUPERT (1999): “Employment and Statistical Discrimination: a Hands-On Experiment,” *Journal of Economics*, 25:1, 85-102.
- [2] ANDERSON, L.R. AND C.A. HOLT (1997): “Information Cascades in the Laboratory,” *American Economic Review*, 87:5, 847-862.
- [3] ANDERSON, L.R., R.G. FRYER AND C.A. HOLT (2002): “Experimental Studies of Discrimination,” in *Handbook on Economics of Discrimination*, ed. by W. Rogers, forthcoming
- [4] ANDERSON, S.P., J.K. GOEREE AND C.A. HOLT (1998): “Rent Seeking with Bounded Rationality: An Analysis of the All Pay Auction,” *Journal of Political Economy*, 106:4, 828-853
- [5] ARROW, K.J. (1973): “The Theory of Discrimination,” in *Discrimination in Labor Markets* ed. by O. Ashenfelter and A. Rees. Princeton: Princeton University Press, 3-33.

- [6] BANERJEE, A.V. (1992): "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, 107:3, 797-817.
- [7] BAYE, M.R., D. KOVENOCK AND C.G. DE VRIERS (1996): "The All-Pay Auction with Complete Information," *Economic Theory*, 8, 291-305.
- [8] BICKCHANDANI, S., D. HIRSHLEIFER AND I. WELCH (1992): "A Theory of Fads, Fashion, Custom and Cultural Change as Informational Cascades," *Journal of Political Economy*, 100:5, 992-1026.
- [9] BULL, C., A. SCHOTTER AND K. WEIGELT (1987): "Tournaments and Piece Rates: An Experimental Study," *Journal of Political Economy*, 95:1, 1-33.
- [10] DAVIS, D.D. (1987). "Maximal Quality Selection and Discrimination in Employment," *Journal of Economic Behavior and Organization*, 8, 97-112.
- [11] DAVIS, D.D. AND C.A. HOLT (1993): *Experimental Economics*. Princeton: Princeton University Press.
- [12] DAVIS, D.D. AND R.J. REILLY (1994): "Do too Many Cooks Always Spoil the Stew? An Experimental Analysis of Rent-Seeking and the Role of a Strategic Buyer," *Public Choice*, 95, 89-115.
- [13] DUFFY, J. AND J. OCHS (2003): "Cooperative Behavior and the Frequency of Social Interaction," mimeo, University of Pittsburgh.
- [14] FILIPPIN, A. (2003): "Discrimination and Workers' Expectations," mimeo, University of Milan.
- [15] FILIPPIN, A. AND A. ICHINO (2003): "Gender Wage Gap in Expectations and Realizations," mimeo, European University Institute.
- [16] FERSHTMAN, C. AND U. GNEEZY (2001): "Some Experimental Evidence on the Evolution of Discrimination, Cooperation and Perceptions of Fairness," *Economic Journal*, 481, 679-703.
- [17] FISCHBACHER, U. (2002): "zTree Experimenter's Manual," mimeo, University of Zurich.
- [18] FRYER, R.G., J.K. GOEREE AND C.A. HOLT (2002): "An Experimental Test of Statistical Discrimination" mimeo, The University of Virginia.
- [19] HARGREAVES-HEAP, S. AND Y. VAROUFAKIS (2002): "Some Experimental Evidence on the Evolution of Discrimination, Cooperation and Perceptions of Fairness," *Economic Journal*, 481, 679-703.
- [20] HEY, J. AND A. MORONE (2002): "Do Markets Drive out Lemmings or Vice-Versa?," mimeo.

- [21] LAZEAR, E.P AND S. ROSEN (1981): "Rank-Order Tournaments as Optimum Labor Contracts," *Journal of Political Economy*, 89:5, 841-64.
- [22] LOHM, S. (2000) "Stereotypes, Sex Discrimination and Paranoia - an experimental study," *Scandinavian WP in Economics* 2000:1.
- [23] MARIMON, R., S. SPEAR AND S. SUNDER (1993). "Expectationally-Driven Market Volatility: Experimental Evidence," *Journal of Economic Theory*, 61, 74-103.
- [24] MURNINGHAM, J.K. (2001): "A Very Extreme Case of the Dollar Auction," mimeo, Kellogg Graduate School of Management, Northwestern University.
- [25] O'KEEFFE, M., W. K. VISCUSI, AND R. J. ZECKHAUSER (1984): "Economic Contests: Comparative Reward Schemes," *Journal of Labor Economics*, 2:1, 27-56.
- [26] PHELPS, E.S. (1972): "The Statistical Theory of Racism and Sexism," *American Economic Review*, 62:4, 659-661.
- [27] ROTH, A.E. (1994). "Let's Keep the Con out of Experimental Econ.: A Methodological Note," *Empirical Economics*, 19, 99-109.
- [28] SCHOTTER, A. AND K. WEIGELT (1992): "Asymmetric Tournaments, Equal Opportunity Laws, and Affirmative Action: Some Experimental Results," *Quarterly Journal of Economics*, 107, 511-39.
- [29] SELTEN, R. (1967): "Die Strategiemethode zur Erforschung des eingeschränkt rationalen Verhaltens im Rahmen eines Oligopolexperiments," in *Beiträge zur Experimentellen Wirtschaftsforschung*, ed. by H. Sauermann. Tübingen: J.C.B. Mohr, pp. 136-168.
- [30] SHUBIK, M. (1971): "The Dollar Auction Game: A Paradox in Non-Cooperative Behavior and Escalation," *Journal of Conflict Resolution*, 15:109-111.
- [31] WOODFORD, M. (1990). "Learning to Believe in Sunspots," *Econometrica*, 58:2, 277-308.

IZA Discussion Papers

No.	Author(s)	Title	Area	Date
810	J. H. Abbring G. J. van den Berg	A Simple Procedure for the Evaluation of Treatment Effects on Duration Variables	6	06/03
811	M. Corak W.-H. Chen	Firms, Industries, and Unemployment Insurance: An Analysis Using Employer-Employee Data	1	06/03
812	J. T. Addison T. Schank C. Schnabel J. Wagner	German Works Councils in the Production Process	3	06/03
813	E. P. Lazear	Firm-Specific Human Capital: A Skill-Weights Approach	5	06/03
814	G. Ridder G. J. van den Berg	Measuring Labor Market Frictions: A Cross-Country Comparison	6	07/03
815	A. Aakvik K. G. Salvanes K. Vaage	Measuring Heterogeneity in the Returns to Education in Norway Using Educational Reforms	6	07/03
816	T. T. Herbertsson J. M. Orszag	The Early Retirement Burden: Assessing the Costs of the Continued Prevalence of Early Retirement in OECD Countries	3	07/03
817	T. M. Andersen T. T. Herbertsson	Measuring Globalization	2	07/03
818	J. Pencavel	The Surprising Retreat of Union Britain	3	07/03
819	M. Beine F. Docquier H. Rapoport	Brain Drain and LDCs' Growth: Winners and Losers	1	07/03
820	C. M. Cornwell K. H. Lee D. B. Mustard	The Effects of Merit-Based Financial Aid on Course Enrollment, Withdrawal and Completion in College	6	07/03
821	P. Carneiro J. J. Heckman	Human Capital Policy	6	07/03
822	D. Weichselbaumer R. Winter-Ebmer	The Effects of Competition and Equal Treatment Laws on the Gender Wage Differential	6	07/03
823	A. Filippin	Discrimination and Workers' Expectations	5	07/03
824	A. Filippin	Discrimination and Workers' Expectations: Experimental Evidence	5	07/03

An updated list of IZA Discussion Papers is available on the center's homepage www.iza.org.