

Filippin, Antonio; Guala, Francesco

Working Paper

Costless discrimination and unequal achievements in a labour market experiment

IZA Discussion Papers, No. 6187

Provided in Cooperation with:

IZA – Institute of Labor Economics

Suggested Citation: Filippin, Antonio; Guala, Francesco (2011) : Costless discrimination and unequal achievements in a labour market experiment, IZA Discussion Papers, No. 6187, Institute for the Study of Labor (IZA), Bonn,
<https://nbn-resolving.de/urn:nbn:de:101:1-201201103667>

This Version is available at:

<https://hdl.handle.net/10419/58581>

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. 6187

Costless Discrimination and Unequal Achievements in a Labour Market Experiment

Antonio Filippin
Francesco Guala

December 2011

Costless Discrimination and Unequal Achievements in a Labour Market Experiment

Antonio Filippin

*University of Milan
and IZA*

Francesco Guala

University of Milan

Discussion Paper No. 6187
December 2011

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0

Fax: +49-228-3894-180

E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series 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 organization supported by Deutsche Post Foundation. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. 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.

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 directly from the author.

ABSTRACT

Costless Discrimination and Unequal Achievements in a Labour Market Experiment^{*}

We investigate the emergence of discrimination in an experiment where individuals affiliated to different groups compete for a monetary prize, submitting independent bids to an auctioneer. The auctioneer receives perfect information about the bids (i.e. there is no statistical discrimination), and she has no monetary incentive to favour the members of her own group (the bidders are symmetric). We observe nonetheless some discrimination by auctioneers, who tend to assign the prize more frequently to a member of their own group when two or more players put forward the highest bid. Out-group bidders react to this bias and reduce significantly their bids, causing an average decay of their earnings throughout the game, with cumulative effects that generate strongly unequal outcomes. Because the initial bias is costless, such mechanism can survive even in competitive market, providing a rationale for a well-known puzzle in the literature, i.e. the long-run persistence of discrimination.

JEL Classification: J71, D44, C9

Keywords: discrimination, tournament, groups, experiment

Corresponding author:

Antonio Filippin
University of Milan
Department of Economics and Business
via Conservatorio, 7
20122 Milan
Italy
E-mail: antonio.filippin@unimi.it

^{*} Research for this paper was financed by a PUR grant and a “Rientro dei Cervelli” scholarship. We are grateful to the experimental laboratory of the University of Parma for logistic support. We would like to thank Olivier Armantier, Anat Bracha, and John Kagel, for helpful suggestions. We also thank the participants at the “6th Nordic Conference on Behavioral and Experimental Economics” (Lund), the conference “Experimental Methods and Economic Modelling” (Capua), the 2011 meetings of the ESA (Chicago) and AIEL (Milan), and the seminar held at the University of Milan for their comments. The usual disclaimers apply.

1. Introduction

Discrimination is a despicable phenomenon that affects in different forms various aspects of social life. It may concern gender, race, social class, geographic origins, ethnicity, age, and several other social categories. For this reason, discrimination has been studied by different disciplines using different methods, and different theories have been proposed to explain its emergence and persistence through time. A prominent explanation, popular among psychologists and biologists, is that discrimination reflects a natural tendency of human beings to favour the members of one's group. This tendency – and the related propensity to penalise the members of other groups – may have been adaptive in our ancestral past, when small groups competed for limited resources, but is maladaptive in the context of large and diverse societies such as those we presently live in.

Explanations based on innate psychological propensities may be translated in the language of economics by introducing an exogenous “preference for discrimination” that influences our decisions in some economically relevant circumstance. One example that has attracted the attention of economists is the labour market, where minority workers may be assigned more menial jobs and/or lower wages even though they have the same skills and productivity as other workers. Traditionally, however, economists have found preference-based explanations unsatisfactory, because discriminatory tastes impose costs on employers that in the long run should be eliminated by market competition. For this reason economists have explored alternative models where discrimination emerges from the interaction of agents who do not have strong discriminatory tastes, but form discriminatory beliefs about workers.

To achieve this result, some theoretical models have introduced uncertainty about the quality of individual workers, while others have focused on cost asymmetries that affect investment in human capital. Some of these models have been tested in the laboratory, where the worker-employer interaction takes the form of a tournament awarding a known prize to the highest bidder in an all-pay auction. A common feature of these experiments is that discrimination is induced *exogenously*, that is, by manipulating the quality of bidders or introducing noise in the information that auctioneers use to discriminate high- from low-quality bidders. However, as we shall

see, the experiments have provided mixed results and only limited support to the models they were meant to test.

In this paper we follow a different route. We describe the emergence of discrimination in an experimental setting with *symmetric bidders* and *perfect information* about the quality of their bids. Following an established tradition in social psychology, we induce discriminatory behaviour imposing an arbitrary *group identity* on bidders and auctioneers. The mechanism that sustains discrimination in this environment is bidders' *expectations*, which are constantly updated following auctioneers' decisions to award the prize in a series of independent tournaments. Expectations are mainly driven by a bias in auctioneers' decisions which is *costless*, and thus cannot be exploited by competitors in a market. This mechanism is intriguing and corrects some flaws of existing models and experiments on the emergence of discrimination. It also bridges the gap between the literature on discrimination in economics, and research on social categorization and group identity in psychology.

The paper is organized as follows: in section 2 we review the existing literature and lay out the problem to be investigated. Section 3 describes the experimental set up and provides a thorough analysis of the data. Section 4 concludes summarizing the main results and articulating their significance for the study of discrimination.

2. Social groups and discrimination

Categorization is a pervasive aspect of human social behaviour, which facilitates information transmission and economizes on our limited cognitive abilities. However, it is generally acknowledged that categorization may also foster discrimination, if our perceptions of the true characteristics of an individual are distorted by stereotypes about the group she belongs to (Campbell 1967, Tajfel 1982, McGarty et al. 2002). This phenomenon is amply documented in the case of minorities such as African-Americans in the United States or Southerners in Italy, who have been and still are burdened with negative stereotypes like idleness, shirking, dishonesty, or lying. Another case in point is gender discrimination, which prevents talented women from gaining positions of leadership in politics, business, and even the academia because of alleged differences in cognitive skills, physical strength, or competitiveness.

The possible causes of discrimination are multifarious, but for theoretical purposes it may be useful to divide them in two broad categories. On the one hand,

discrimination may reflect an underlying *preference* to treat more favourably the members of one's own group than the members of a group which is perceived to differ along some dimension. On the other hand, discrimination may be a consequence of people's *beliefs* which associate (rightly or wrongly) a certain set of characteristics with membership in a certain group. These two mechanisms (preference-based and belief-based discrimination) are not mutually exclusive, of course, and may play different roles in different contexts. In this paper we use them to organize our survey of the literature in separate sub-sections.

2.1. Preference-based explanations of discrimination

The simplest explanation of discrimination is that people like it. Since the 1960s psychologists have gathered abundant evidence that human beings of all cultures and faiths have a strong tendency to treat preferentially the members of their own group compared to the members of other groups. In a famous field experiment with middle-class teenagers, Muzafer Sherif showed that the mere creation of group identities in the context of a peaceful summer camp increased significantly the level of competition and aggressiveness (to the point that the experiment had to be suspended – see Sherif et al. 1961). Henri Tajfel and his collaborators (e.g. Tajfel et al. 1971, Tajfel 1982) pursued a similar line of research using the so-called “minimal group paradigm”, an experimental setting where subjects' behaviour is manipulated creating artificial groups based on arbitrary criteria and meaningless labels. Again, it turns out that individuals become more cooperative, altruistic and caring towards the members of their own group than the members of other groups. From an economic point of view, it is noteworthy that group members are willing to pay a cost to increase inter-group differences in earnings or achievements (they are inefficiently spiteful towards out-group members, in other words).

In-group favouritism is not entirely mysterious from an evolutionary point of view (Richerson and Boyd 2001). In the ancestral past cooperation must have provided *homo sapiens* with a comparative advantage in fitness terms. The costs of altruism and cooperation were probably recouped in the course of repeated encounters with the members of one's own family or tribe, while competition for territory and resources made suspicion and hostility advantageous in the context of inter-group behaviour (e.g. Choi and Bowles 2007). The same logic however implies that the suppression of discrimination *within* the group should be advantageous, for a larger cooperative group

should be more efficient and hence out-compete smaller groups. In larger, diverse societies like ours then there may be considerable efficiency losses caused by our innate tendency to favour individuals that we perceive as more similar in terms of race, ethnicity, etc.

This “paradox of discrimination” (discrimination exists even though it is inefficient) emerges in a similar guise in economic theories that posit an exogenous preference for discrimination. In a seminal series of papers Gary Becker (1957) proposed an explanation of discrimination in the labour market that is driven entirely by the preferences of employers, customers or co-workers. While the generality of this model makes it applicable to a number of real-world cases (sex, religion, race, etc.) it also raises the puzzle of the persistence of discrimination in a competitive environment. If discrimination is costly, it should be wiped away by market competition in the long run, in spite of abundant evidence to the contrary. So partly out of dissatisfaction with preference-based models, economists started to devise in the 1970s models that can explain the existence of discrimination equilibria in competitive markets. These models typically posit an invisible-hand mechanism that generates discrimination independently of peoples’ tastes for discrimination, and are briefly reviewed in the next section.¹

2.2. *Belief-based explanations of discrimination*

If not by their preferences, discriminatory behaviour may be generated by employers’ *beliefs* in the inferior quality of workers belonging to a certain group. For discrimination to persist over time, however, it is necessary that such beliefs reflect a genuine difference between workers’ productivity – or, to put it differently, that employers do not have the opportunity to learn the true quality of discriminated workers. Since the 1950s social scientists have been aware that in unhappy circumstances discriminatory beliefs may become a *self-fulfilling prophecy* (Merton 1957), if workers have an incentive to conform to employers’ expectations. This outcome is not only unfair but also socially inefficient, of course, because lower levels of effort or investment in human capital by a sizeable minority cause losses of productivity at the aggregate level.

Arrow (1973) modelled discrimination as a self-fulfilling prophecy in a seminal paper devoted to discrimination in the labour market. Arrow’s model assumed

¹ For more comprehensive surveys, see also Rodgers (ed. 2006), and in particular the chapter on experiments by Anderson et al. (2006).

incomplete information regarding the quality of workers, who must decide how much to invest in training before they enter the job market, and condition their investment decision on employers' earlier behaviour. Arrow assumed that employers enter the market with asymmetric prior beliefs concerning the distribution of human capital across two classes (A and B) of workers. Given the same signal, therefore, an employer will prefer to hire more skilled A-workers rather than inferior B-workers. This will provide an incentive for B-workers to invest less in training, which will be taken by employers as further confirmation of their inferior quality. And so forth: the prior belief has become self-fulfilling.

Arrow started the branch of so-called theories of "statistical discrimination" (for a thorough survey see Fang and Moro 2010). In such theories, typically, ex-ante equal groups achieve unequal outcomes ex-post. Although it is not individually costly in the sense of taste-based models, statistical discrimination has been challenged for relying on the assumption that employers cannot discover workers' true skills by means of trial work periods (see Aigner and Cain 1977, Cain 1986). In a tournament for example the period before the promotion decision can be considered a trial period in which workers' characteristics can be observed rather than inferred. However, as shown by Filippin (2009), minority workers who expect to be discriminated against face a lower expected return on their effort and may behave differently even in the trial work period. As a result, even though in equilibrium there may be profit opportunities for firms that are willing to pay B-workers more, no firm would be able to see them.

This explanatory advantage of Arrow's model however comes at a price. An obvious question concerns the origins of the prior beliefs that trigger the self-fulfilling prophecy. In an early experiment, Davis (1987) conjectured that biased prior beliefs may result from a simple statistical-cognitive distortion: the sample of candidates from the minority group, being smaller, is likely to contain fewer top-quality candidates. If employers remember top quality candidates more vividly, they will form the impression that on average the quality of minority workers is lower. Another possibility is that the priors reflect previous experience. Fryer, Goeree and Holt (2005) describe a classroom experiment where employers observe a noisy investment signal and discover the real quality of workers only after they have made their hiring decision. For a few rounds, employers sample from two groups of workers (Green and Purple) with asymmetric quality. When the asymmetry is removed, employers' hiring decisions are still based on the (wrong) expectations that Green workers are better than Purple ones.

As it turns out, however, neither mechanism seems powerful enough to generate robust discrimination. In Davis' (1987) experiment, the sampling bias generates only mild discrimination, unless it is implemented with a particularly heavy hand. Similarly, in Fryer et al. (2005) the self-fulfilling prophecy is quite fragile: when the asymmetry is removed, the Purple (previously disadvantaged) players see new opportunities and raise their bids substantially, causing a surge in their hiring rate. Employers thus learn quickly about the new (symmetric) distribution of quality, and discrimination disappears.

A similar phenomenon has been observed by Filippin (2008) in an experiment with human bidders (workers) facing an artificial auctioneer (employer). The auctioneer is programmed to hold discriminatory preferences against one group of bidders during the early rounds ("crazy" auctioneer condition), and to become a "fair auctioneer" at a later stage of the game. Contrary to the Fryer et al. (2005) setting, in Filippin's experiment the bidders are not informed of the change that occurs in the parameters of the game – which intuitively should make the persistence of statistical discrimination more likely. Nevertheless, in most sessions minority bidders learn about the new (more favourable) environment by occasionally submitting high bids that turn out to be successful. So discrimination seems to be less resilient in this laboratory environment than one would expect from statistical discrimination theory.

Another anomalous finding that may interfere with statistical discrimination is the propensity of disadvantaged players to overbid in asymmetric tournaments, reported by Bull, Schotter and Weigelt (1987). This behaviour may be motivated by a desire to compensate what subjects perceive as an "unfair" disadvantage. However, it may prevent statistical discrimination from occurring, by effectively "levelling up" the exogenous inequalities introduced by experimenters.²

So far, to conclude, there has been only limited experimental research on the emergence of discrimination in tournaments. The existing evidence is mixed, and does not indicate strongly one mechanism that is responsible for the emergence and persistence of discrimination. Belief-based mechanisms such as Arrow-style statistical discrimination are empirically fragile, while preference-based explanations should be

² It is also possible to level the field by policy intervention, for example introducing affirmative action or equal opportunity programmes. The effects of such policies in experimental tournaments have been studied by Schotter and Weigelt (1992).

vulnerable to competition in markets. Jointly, these considerations suggest that there may be alternative ways to explain the emergence and resilience of discrimination.

In the rest of the paper we describe an experiment that points in an interesting new direction. The experiment lies at the crossroad between the literature on group behaviour and the literature on tournaments. As in classic group identity experiments, subjects are divided in two arbitrary groups. Experimental subjects then participate in a series of tournaments with feedback at every round. The tournaments are perfectly symmetric – that is, we do not impose different investment costs on bidders. Whatever discrimination will emerge then is certainly due to the well-documented effect of group identity on players' preferences. Bidders observe auctioneers' decisions, and update their discrimination beliefs. We should then expect some equilibrium between auctioneers' preferences and bidders' expectations to emerge with repetition, and the nature of such equilibrium may provide valuable insights in the mechanics of discrimination. In the next section we describe this setting in more detail, describe the results, and outline an explanation based on a Bayesian model of learning.

3. The experiment

In a *tournament* the payoff of every agent depends on her relative performance, i.e. how well she has done compared to other participants. A common way to determine payoffs in a tournament, for example, is to rank participants based on performance and to assign *prizes* in a descending order (the highest prize to the best performing agent, the second-highest prize to the second agent, and so on). Tournaments can be used to model a number of familiar situations in real life, such as political elections, promotions within a firm, or competition for a new job. In the latter case, the prize is the job's salary and applicants compete submitting their CVs and performing during interviews. Their chances of being hired, however, typically depend on previous investments in education, training, etc. Because these investments are expensive, the costs may be represented as bids that are deducted from the prize. In other words, the total payoff is given by the prize minus the cost of the bid. It is generally assumed that a higher bid buys a higher chance to win the prize, because the employer prefers to hire more qualified applicants with higher human capital. The employer then acts as an auctioneer who observes the bids (investments) of the various applicants, chooses a bid that maximizes her utility, and assigns the prize to the respective bidder. The other applicants, who do not receive

the prize, suffer a cost that is equal to the investment they have made: the tournament takes the form of all-pay auction in which all bidders pay their bids regardless of whether they have won the prize or not.

3.1. Experimental setting

We implemented this setting experimentally letting each subject play a series of ten tournaments with identical parameters, in groups of five players (four bidders and an auctioneer). To eliminate the potentially confounding effect of reputation, we used a “stranger” design where every subject was randomly allocated in a new group at every round. The role of each subject (bidder or auctioneer) was randomly determined at the start and remained the same throughout the experiment.

At each round, each bidder received an endowment of 60 cents of a euro that she could invest to win a prize of 400 cents awarded by the auctioneer to a single bidder. Bidders were perfectly symmetric – we did not impose different investment costs – and their payoffs in each round were equal to their endowment minus their bid plus the prize (if they won). The endowment could not be accumulated over a series of rounds: after each round, whatever fraction of the endowment had not been spent was automatically paid and a new endowment of 60 cents was provided for the next round. The auctioneer’s payoff in contrast amounted to the winning bid she had chosen, multiplied by a factor of three (to increase the difference between bids, in case it was small).

The monetary payoffs were designed in such a way as to eliminate any rational incentive to underbid. Suppose in fact that the auctioneer is unbiased (i.e. she chooses one of the highest bids that have been submitted and is indifferent between equal bids). If every bidder invests 60 cents (their whole endowment), the expected monetary value of her bid is

$$EV(\text{bid}60) = (0 \times .75) + (400 \times .25) = 100.$$

Since the expected monetary value of keeping her endowment is 60 cents, a rational risk-neutral player should always bid her whole endowment, assuming that all the opponents are doing the same. An unbiased auctioneer has no incentive to deviate from a random decision and this strategy profile is therefore part of an equilibrium. Assuming a profit-maximizing (i.e. unbiased) auctioneer who randomizes to break the ties, such an equilibrium is also unique, since bidders are always strictly better off if they bid 60 (either they outbid the opponents, or they match their proposal).

This reasoning depends crucially on the assumption of auctioneer neutrality: the probability of winning the prize clearly declines if the auctioneer discriminates against some bidders. To introduce this possibility we added a preliminary phase to the experiment in which an arbitrary “group identity” was primed using standard techniques borrowed from the experimental literature in social psychology.

Social psychologists have identified several factors that contribute to the creation of “group identity”, including physical proximity, face-to-face contact, perceptual similarity, interdependence, and common fate (Hogg and Abrams 2003). As they entered the laboratory – and before they participated in the tournament – twenty subjects were randomly assigned to two groups (named “Red” or “Blue”) and asked to wear a bracelet of the corresponding colour. Each group then performed a simple cooperative task aimed at reinforcing their group identity: the goal was to memorize at least ten verses of a short poem in no more than five minutes, with another five minutes to write them down. The two groups performed the task in separate areas of the laboratory, under the supervision of an experimental assistant who did not interfere with their work but simply watched the time and monitored the results.

The task was designed to suggest a natural division of labour among the group members (each member could memorize one verse), and to almost guarantee the success of every group. Successful completion was rewarded with a prize of 30 euro, to be equally divided among the ten members of the group.³ These two factors (interdependence and common fate) magnify the bias induced by coloured bracelets, physical proximity, and face-to-face contact.

At the end of this preliminary phase all players were asked to sit at their randomly assigned computer terminals, divided by partitions. An assistant read the instructions aloud, while the subjects followed on paper. Questions from the audience were elicited, until all misunderstandings had been resolved. At this point the second phase of the experiment (the tournaments) began. The first screenshot provided information concerning the role of each subject in the experiment: sixteen participants played as bidders (eight Blue, eight Red), and four played as auctioneers (two Blue, two Red).

³ All fourteen groups who participated in the experiment successfully completed the task, which they seemed to enjoy.

Two Blue and two Red bidders participated in each tournament together with a Red or Blue auctioneer, with reshuffling at every round. The colour of each bidder was either the same or different from the auctioneer's colour. Following an established terminology in the experimental literature, we call *in-group* members those bidders who shared the same colour with the auctioneer, and *out-group* members those who did not. The in-group/out-group allocation remained fixed throughout the experiment. In other words: even though the auctioneer could change at every round, her colour remained the same. Each bidder was either always in-group, or always out-group. Auctioneers faced two in-group bidders and two out-group bidders in every round.

At all stages of the experiment the identity of players was kept secret, while their group affiliation was common knowledge. The division in Red/Blue groups was meant to elicit group identity and, possibly, group favouritism in the tournaments. Of course there was no a priori guarantee that group priming would have a significant effect on subjects' behaviour, but as we shall see the data strongly confirm this hypothesis. Our design then differs from all previous experiments where discrimination was induced imposing different investment costs on bidders. In contrast we manipulated directly auctioneers' attitude towards discrimination, and saw whether this manipulation had an effect on their behaviour and the behaviour of bidders. To make this setting particularly inhospitable to discrimination we also eliminated all sources of uncertainty, such as noisy signals or asymmetric knowledge of payoffs. In our experiment all monetary payoffs were common knowledge, and there could not be statistical discrimination – i.e. the investments (bids) were perfectly observed by auctioneers. Complete feedback about submitted bids, as well as the choice of the auctioneer, was provided after every round. So the only source of incomplete information was the possible discriminatory bias of the auctioneers, about which bidders had to form (and update) beliefs.

Notice that expectations of discrimination change the expected value of bids, and may deter competition in the tournaments. Suppose for example that a fraction d of auctioneers discriminates against out-group players. The expected value of bidding the full endowment, given that one is an out-group bidder, then becomes

$$EV(\text{bid}60|\text{out}) = d \times 0 + (1 - d) [(0 \times .75) + (400 \times .25)] = 100 - 100d.$$

Given that $EV(\text{bid}0) = 60$, a risk-neutral out-group player should bid nothing if she believes that $d > 0.4$. Assuming that the auctioneer actually discriminates against out-groups, such a behaviour is also part of an equilibrium. Notice that auctioneers can discriminate against out-group bidders without forfeiting any profits, because they can

do it at no cost. The reason is that in-group bidders have an incentive to compete among themselves anyway, even if they think that out-group players have no chance of being awarded the prize. In equilibrium in-group players should always bid 60, because lower bids would grant the same-colour opponent a profitable opportunity to win by bidding the maximum amount (assuming that the auctioneer cares about monetary payoffs).

Consider, moreover, that bids in the interval [1-59] can never be rationalized for either type of player. Assuming that bidders are (and they believe that their opponents are) sufficiently sophisticated iterative reasoners, and that the auctioneer cares about monetary payoffs, the optimal bid is either zero or 60. Even relaxing the strong assumption that players play the equilibrium strategies via introspection, they should learn that bids in the interval [1-59] cannot be an equilibrium by means of trial and error (evolutionary) learning. Therefore, we expect bids in the interval [1-59] to disappear with the repetition of the game.

To sum up, our experimental setting is characterized by multiple equilibria that depend on out-group bidders' and auctioneers' behaviour. In principle, we could observe:

1. symmetric bids and unbiased auctioneers;
2. asymmetric bids and discriminatory auctioneers.

In these two cases out-group beliefs would be correct both at the beginning of the game, and in equilibrium. However, there may also be self-confirming equilibria driven by false prior beliefs, such as

3. asymmetric bids driven by *wrong* prior expectations of discrimination (because the auctioneers are unbiased).

Case 3 may trigger an Arrow-style self-fulfilling mechanism, where incorrect priors lead to underinvestment, and this in turn generates unequal outcomes. Case 2 is similar, except that bidders' prior beliefs are not wrong but reflect a genuine bias on auctioneers' part. Case 1 is in a sense the least interesting for it would mean that we have failed to implement discrimination in the laboratory.

Which one of these equilibria is instantiated may depend of course on the dynamic of the game. Because beliefs may evolve during the experimental game, we shall focus both on bidders' "naïve" strategies at the beginning of the game, and on their behaviour at the end of the experiment, after they have received feedback concerning the decisions of auctioneers. To summarize, we shall ask the following questions:

QUESTION 1: Do out-group players bid more/less than in-group players at the beginning of the game?

QUESTION 2: Does the behaviour of in-group or out-group players change significantly during the experiment?

QUESTION 3: Do out-group players bid more/less than in-group players at the end of the game (in equilibrium)?

Another set of questions concerns the factors that determine behaviour in equilibrium. In particular:

QUESTION 4: Are auctioneers biased in favour of in-group members, and are they willing to pay a cost to discriminate against out-group bidders?

QUESTION 5: Is the behaviour of bidders influenced significantly by auctioneers' decisions, or do prior beliefs largely determine the dynamics of the game?

We try to answer these questions in the next section, where we describe the experimental data and provide a statistical analysis of the main findings.

3.2. Experimental results

We report the data of seven sessions with 140 subjects in total which we ran at the University of Parma, Italy, over two days in April and May 2011. Group allocation was designed in such a way as to obtain two independent observations per session, or fourteen overall. Most subjects were in their low twenties, and studied business and economics. Average earnings were approximately 16 euro, for about one hour of experimentation. Auctioneers earned on average more than bidders (20 vs. 15 euro) but the greatest achievers in absolute terms were a small subset successful bidders, who took home as much as 28 euro. The experiment was designed according to the usual conventions of experimental economics.

We begin our analysis with a general survey of auctioneers' and bidders' behaviour. Figure 1 includes the average bidding patterns of in-group and out-group bidders for each session of the experiment (*1b-h*) as well as aggregate data from all sessions (*1a*). It is immediately apparent that in most sessions in-group players tended to bid more than out-group players (session 7 was the exception). It is also apparent that this difference emerged early, but not immediately in the experiment: in the first round

there was a prevalence of out-group bids. Once established, however, the in-group out-group difference was fairly robust throughout the game: in equilibrium, in-group players bade more.

[Figure 1a-h about here]

We test these propositions using data from all sessions. At the beginning of the game in-group players bade on average 33.1 cents in the first round, compared with 41.0 cents offered by out-group players. Since there has been no strategic interaction among players before the first bid, we regard all the 112 bids as independent observations. The difference is statistically significant ($P > |z| = .046$).

RESULT 1: *Out-group players bid more than in-group players in the first round of the game.*

This asymmetry however was reversed already in the second round, and the opposite trend persisted until the end of the experiment. The average bid throughout the ten periods was 43.1 cents for in-group players and 34.2 cents for out-group players, a statistically significant difference at $P > |z| = .005$.

RESULT 2: *In-group players bid more than out-group players throughout the game.*

These two results are interesting because they suggest that the group identity manipulation administered in the first phase of the experiment did not discourage out-group bidders at the start. Strictly speaking, it is impossible to rule out the existence of expectations of discrimination in round one. Out-group players' rather bullish bidding in fact may be interpreted as an attempt to compensate a perceived disadvantage – a phenomenon observed by Bull, Schotter and Weigelt (1987) in a different context. Behaviour in the first round however should be taken with a pinch of salt, for it certainly reflects a certain amount of confusion and mistaken beliefs. One sign of confusion are intermediate bids (bids between zero and sixty cents), which should not be observed if players are sufficiently sophisticated iterative reasoners. Table 1 summarizes the percentage of maximum (60), minimum (0) and intermediate bids ($0 < \text{bid} < 60$) submitted by in-group and out-group players across the ten rounds of the

experiment. We observe in both groups a decline of intermediate bids, from 60-70% in period 1 to about 20% in period 10, which suggests that a process of learning has taken place during the course of the experiment. (It is noteworthy that the sharpest reduction in intermediate bids occurred after the first period, and that learning was probably over by period six.)

[Table 1 about here]

In light of these data, and following a useful convention in experimental economics, we consider behaviour at the end of the game as most indicative of players' preferences and beliefs in equilibrium. We take data from the last two rounds. Here the average bids amount to 42.4 for in-group players and 33.3 for out-group players, a difference that is statistically significant at $P > |z| = .08$.

RESULT 3: In-group players bid more than out-group players at the end of the game (in equilibrium).

Table 1 shows that at the end of the game 64% of in-group players bade the maximum amount (60 cents), compared with only 46% of out-group players. This difference however was already manifest in period three, and persisted with minor variations thereafter. So it is not surprising that in-group players were on average more successful than out-group players. At the end of the experiment, in-group bidders took home on average 17.5 euro, compared with 12.7 euro earned by out-group bidders. Figure 2 summarizes the proportion of winning bids by in- and out-group players during the course of the experiment, aggregating across all sessions. Notice that in-group players won the majority of tournaments in every round except one (round nine). Overall, they won the prize 64.3% of the time. What has to be explained is whether the behaviour of auctioneers was overall rational, in the sense that they tended to maximize their profits by choosing the highest bids, or was instead driven by discriminatory tastes.

[Figure 2 about here]

In order to test this proposition we have to look more carefully at auctioneers' decisions. Table 2 reports the decisions made by auctioneers, given the distribution of

high bids in the tournaments.⁴ There are eight logically possible cases, depending on the number of high bids (one, two, three, or four), and their type (in- or out-group). In the first two columns (“Highest bids made by”) we find the possible combinations of high bids; the third and fourth columns (“Number of winners”) report the decisions of auctioneers for each combination of high bids. For example: the first line includes data from all tournaments where two in-group players submitted the highest bids. This happened 27 times in the experiment (last column), and in 23 cases (third column) the auctioneer awarded the prize to an in-group player (in other words, she maximized her profits 85% of the time).

[Table 2 about here]

It is instructive to compare the top two lines with the bottom two lines of Table 2. These include data about the tournaments in which the high bids were submitted either all by in-group or all by out-group players. Notice that in-group bids were unmatched by out-group bids 81 times, whereas the opposite situation occurred only 39 times. In either circumstance, the auctioneer chose in the overwhelming majority of cases to give the prize to one of the high bids. Only nine times (11.1%) a lower out-group bid was preferred to a higher in-group bid; the opposite happened four times (10.2%). Notice that while choosing a lower in-group bid may be considered a case of *costly* discrimination, choosing a lower out-group bid is a case of “*reverse*” costly discrimination. Given the low frequency and symmetry of these anomalous data, however, they should be considered random variations in an otherwise systematically rational pattern of auctioneer behaviour.

Lines 3 to 6 in Table 2 (shaded) include data about tournaments where high bids were submitted by in-group *and* out-group players at the same time. We shall refer to these cases as *tie-breaks*. Tie-breaks are an important element of our experiment, for they provide auctioneers with an opportunity to engage in *costless discrimination*. In tie-breaks auctioneers may send a discriminatory signal while still maximizing their profit. This is not an unrealistic event, given that in many real circumstances an

⁴ “High bid” from now on means the highest submitted bid in a given round. Since the bidders are symmetric (they have the same endowment), there must be at least one and there cannot be more than four high bids per tournament.

employer is likely to face several equally qualified applicants. Notice that in such a case it would be inappropriate to speak of a *preference* for discrimination, because auctioneers would not be willing to pay a monetary cost to engage in favouritism. To distinguish, we shall speak of a discriminatory *bias* in tie-breaks, and leave open the question whether the bias is rooted in a cognitive heuristic (in-group bids might be more salient, for example) or reflects an asymmetry in auctioneers' motives (a desire to benefit in-group players).

In tie-breaks there was an obvious asymmetry in the choices of auctioneers, with 65% of the prizes going to in-group members. The null hypothesis that such a distribution of outcomes derives from a binomial distribution that mimics an unbiased auctioneer can safely be rejected ($P > |z| < .0001$). The number of high bids submitted by each type of player does not affect the result significantly: when auctioneers were choosing between an *equal number* of in-group and out-group high bids, in-group players won the tournament 66% of the time. When two high bids out of three were submitted by in-group members, their probability of winning was 68%. When two out of three high bids were submitted by out-group players, in-group players won 52% of the tournaments

RESULT 4: *Auctioneers favour in-group members when discrimination is costless (in tie-breaks).*

Given this result, it is natural to ask whether tie-break discrimination was the main determinant of the decline of out-group contributions. Consider that in- and out-group bids began to diverge already in round two, and were more or less stable by round five. In the first round in-group bidders won almost 60% of the tournaments, even though they submitted on average *lower* bids. Consider however that aggregate data were not available to individual bidders during the game. Each player could only observe the distribution of bids and the outcome of the auctions in which she was directly involved. The relevant question then is how this feedback was processed, and how it influenced bidders' strategies as they collected more evidence of discrimination.

At the beginning of the game bidders did not know whether auctioneers were going to favour in-group members and discriminate against out-group players. They could make conjectures, of course, but given that out-group bidders did not bid less than in-group players in the first round it is reasonable to assume that bidders entered the

game with a prior probability estimate of 0.5, which is equivalent to assuming that there was no initial expectation of discrimination, or perhaps with a mild expectation of discrimination against out-group players, which did not deter their bids significantly. Starting from the first round, bidders received signals regarding auctioneers' discriminatory tastes. Learning can be represented as an updating process of each individual's (estimated) likelihood that the prize will be assigned to a bid submitted by a member of her group, given the meaningful signals that she received as the game proceeded.

A signal is *meaningful* if it provides useful information regarding auctioneers' propensity to discriminate. If all high bids have been submitted by in-group bidders, for example, and the auctioneer assigns the prize to a high bid, the signal is not meaningful according to our definition, since the choices of a biased and of an unbiased auctioneer would coincide. If there is a tie, in contrast, and the auctioneer chooses an in-group bidder, the signal may be interpreted as evidence of discrimination. Similarly, we define the signal as meaningful if the auctioneer awards the prize to a colour that did not make the highest bid, i.e. in case of costly discrimination. In the next section we explain bidders' behaviour using a Bayesian updating model that exploits the objective informative content of each signal.

3.2.1 A Bayesian model of learning

We model the situation that bidders face as a sampling problem from an urn with an unknown proportion of "in" and "out" balls, with replacing at every round. If the fraction of "in" and "out" balls is not the same, the urn represents auctioneers' stochastic propensity to discriminate. Bidders try to learn the true value of that propensity observing a sequence of draws from the urn (auctioneers' decisions). At each round a bidder observes one ball (decision) and updates her probability estimate of the in/out ratio of balls in the urn.

Bidders' beliefs about the composition of the urn can conveniently be represented using a Beta distribution, the shape of which is characterized by two parameters (α , β) that in our case represent the number of positive and negative signals received by a player in during the game. At the beginning of the game, when no signal has been observed yet, we can interpret the parameters of the Beta distribution as already incorporating a number of hypothetical signals equal to $\alpha-1$ and $\beta-1$, respectively. For instance, prior expectations that take the form of a *Beta*(1,1) are

equivalent to assuming an entirely open-minded bidder who assigns a uniform probability over the whole $[0, 1]$ spectrum, as if no hypothetical signal has been already incorporated. The uniform distribution is in fact a special case of the Beta distribution, when both parameters are equal to one (see Figure 3a).

There are of course other possibilities, but as long as the symmetry is preserved ($\alpha = \beta$) higher values correspond to steeper bell-shaped distributions with a mean of 0.5, i.e. they correspond to the beliefs of a bidder who expects to meet an unbiased auctioneer. For instance, a *Beta*(2,2) is represented in Figure 3b. It corresponds to the beliefs of a bidder who expects to face a fair auctioneer, as if she had already observed one positive and one negative signals. Intuitively, the higher the prior value of the two parameters, the stronger her initial confidence that she is facing a fair auctioneer, and the slower the convergence of her beliefs towards the true composition of the urn, as new evidence accumulates. The convergence properties of both the *Beta*(1,1) and the *Beta*(2,2) are depicted in Figure 3c, assuming a sequence of consistent positive signals.

[Figure 3 about here]

The values of α and β do not need to coincide. Different values represent the beliefs of a bidder who expects to face a biased auctioneer. The higher β is relative to α , the more skewed to the right the distribution is, and the higher the expectation of being discriminated against. We initially assume that both in-group and out-group bidders are characterized by a *Beta*(1, 1), i.e. that both populations are characterized by prior expectations of facing an unbiased auctioneer. This assumption is consistent with the data of the first period, but robustness checks will be performed to take into account alternative specifications.

In subsequent rounds, beliefs are updated in a Bayesian manner, given the signals that are actually observed. The characteristics of a Beta combined with a binomial distribution imply that posterior beliefs are also distributed as a Beta with the parameter α (β) increasing one unit at a time as each positive (negative) signal is observed. As explained earlier, we assume that the updating process is triggered only by meaningful signals, i.e. either by choices that imply costly discrimination or by tie-break decisions.

The next step is to estimate the *reaction function* of each type of player (in-group and out-group), defined as the relation between the likelihood that one's colour is going to win and the submitted bid. The goal is to account for players' behaviour given their expectations throughout the whole game, but especially in the early rounds of the experiment where the most drastic adjustment in bidders' behaviour takes place. We estimate the reaction function using the following empirical specification:

$$B_{i,t} = \gamma_0 + \gamma_1 L_{i,t} + \gamma_2 L_{i,t}^2 + \gamma_3 L_{i,t} * OUT + \gamma_4 L_{i,t}^2 * OUT + \varepsilon_{i,t} ,$$

where $B_{i,t}$ is the bid submitted by player i at time t , L is the likelihood that the prize will be assigned to a bidder of one's own type, and the γ s are parameters. The dummy variable OUT is introduced to account for the possibility that the reaction functions of in- and out-group players differ, or in other words that bidders from different groups react differently to the same signals. This is necessary to account for the possible effect of the group identity manipulation we introduced in phase one of the experiment. Moreover, the model accounts for possible non-linearities. Given the nature of dynamic panel of the data, in which the bid in every period depends on the whole history of what happened before, an OLS estimate would return biased and inconsistent estimates. Hence, we estimate it using GMM-type instruments for the difference equation and computing the standard errors with the Arellano-Bond robust VCE estimators.

[Table 3 about here]

Results are presented in the first column of Table 3 and show that the hypothesis of different reaction functions is *not* supported by the data. In fact, the interaction terms are not significant. This suggests that the unequal achievements observed in our experiment are driven by auctioneers' decisions, rather than by prior expectations of discrimination, or by a difference in the reaction functions of in-group and out-group players.

[Figure 4 about here]

The shape of the reaction function implied by the estimated coefficients is represented in Figure 4. The vertical line marks the 50% likelihood (the belief that the auctioneer is unbiased). A second interesting feature of the estimated reaction function

is that both in-group and out-group players reacted asymmetrically to positive and negative information: if they believed that the auctioneer was discriminating against them, they reduced their bids more than they were willing to increase them if they believed that she was biased in their favour.

From a theoretical point of view, the bid of a risk-neutral bidder who expects the opponents to bid the whole endowment jumps from 60 to 0 around a probability equal to .40, as we have seen. However, no such discontinuity can be detected in our data. If we estimate a specification including a cubic term, to allow the model to fit a possible inflection point around the .40 likelihood, all coefficients lose significance. As a robustness check, we also test models with different prior beliefs. Our results do not change qualitatively if we assume that out-group bidders have mild prior expectations of being discriminated against (about 1/3 probability, summarized by a Beta(1,2)), while in-group players believe that the auctioneer is unbiased. Coefficients are reported in the second column of Table 3. More generally, the specification and the results are robust to small perturbations of prior beliefs that go in the right direction (i.e. in-group players are more likely to win), provided the priors are flexible enough to quickly adapt to the evidence that accumulates in each period.⁵ An alternative specification with stickier prior beliefs, for instance as reflected by a Beta(2,2), would be unable to fit the data.

RESULT 5: Assuming Bayesian updating of flexible priors, all players (in-group and out-group) react in the same way to signals, but are more sensitive to negative than to positive signals.

This suggests that the group identity manipulation did not have an effect on the way in which bidders “read” or reacted to the signals sent by auctioneers. The main difference – which explains the decay of their bids – is that out-group players received a greater amount of negative signals early in the game, compared to their in-group fellows. This created negative expectations of discrimination, which led to a reduction in out-group bids. In equilibrium, auctioneers engaged in costless discrimination, out-group

⁵ Of course, this does not hold if beliefs are widely mistaken. For instance, assuming prior beliefs that out-group players are more likely to win would lead to significantly different reaction functions across groups.

players bade less, and as a consequence *achieved* less in spite of perfectly symmetric costs and the absence of statistical discrimination.

4. Conclusion

Discrimination is a familiar and yet elusive phenomenon. Models that postulate preferences for discrimination are vulnerable to a simple objection: such preferences create arbitrage opportunities that should be exploited by non-discriminating firms in a competitive market. Models that postulate incorrect prior beliefs face the problem of justifying the origins of these beliefs, and must explain why workers cannot send signals that correct employers' prior estimates. Such signaling has been observed in laboratory experiments, where statistical discrimination has turned out to be fragile to repetition.

In this paper we have reported data from an experiment with artificially created group identities, where discrimination emerged quickly and persisted through time. The basic setting is a tournament (all-pay auction) where four bidders submit offers to an auctioneer, who has an incentive to assign a monetary prize to the highest bid. Bidders are completely symmetric, and in equilibrium should submit identical bids if the auctioneer is unbiased. Using a classic device from social psychology, we induced a slight bias in auctioneers' decisions. This bias turned out to be strong enough to deter high bids from out-group players, but not enough to create a costly preference for discrimination. The discrimination pattern that we observed, therefore, cannot be eliminated by market competition in the long run.

Further analysis reveals that discrimination is mainly driven by auctioneers' behaviour. There is no evidence that out-group subjects had pessimistic expectations of discrimination at the beginning of the experiment. Nevertheless, large and significant differences between the bids (and outcomes) of the two groups emerged quickly with repetition. This is the effect of two mechanisms: on the one hand, bidders received signals that indicated a mild discriminatory bias in auctioneers' decisions, and updated their expectations accordingly. On the other, bidders strongly reacted to negative signals, and drastically revised their bids downwards. Discrimination, thus, was sustained by players' different reactions to positive/negative signals, combined with the fact that out-group players received more negative signals, rather than by an anomalous reaction of the out-group bidders.

References

Aigner, D.J. and Cain, G.G. (1977): "Statistical Theories of Discrimination in Labor Markets", *Industrial and Labor Relations Review*, 30: 175-87.

Anderson, L.R., Fryer, R.G., and Holt, C.A. (2006) "Discrimination: Experimental Evidence from Psychology and Economics", in *Handbook of the Economics of Discrimination*, ed. by W.M. Rodgers. Cheltenham: Edward Elgar.

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.

Becker, G.S. (1957) *The Economics of Discrimination*. Chicago: University of Chicago Press.

Bull, C., Schotter, A. and Weigelt, K. (1987) "Tournaments and Piece Rates: An Experimental Study", *Journal of Political Economy* 95: 1-33.

Cain, G.G. (1986) "The Economic Analysis of Labor Market Discrimination: a Survey", in *Handbook of Labor Economics*, ed. by O. Ashenfelter and R. Layard. Amsterdam: North Holland, 693-785.

Campbell, D. (1967) "Stereotypes and the Perception of Group Differences", *American Psychologist* 22: 817-829.

Choi, J.K. and Bowles, S. (2007) "The Coevolution of Parochial Altruism and War", *Science* 318: 636-640.

Davis, D.D. (1987) "Maximal Quality Selection and Discrimination in Employment", *Journal of Economic Behavior and Organization* 8: 97-112.

Fang, H. and A. Moro (2010) "Theories of Statistical Discrimination and Affirmative Action: A Survey", NBER Working Paper 15860.

Filippin, A. (2008) "Discrimination and Workers' Expectations: a View from the Lab", mimeo.

Filippin, A. (2009) "Can Workers' Expectations Account for the Persistence of Discrimination?" IZA Discussion Paper 4490.

Fryer, R.G., Goeree, J.K. and Holt, C.A. (2005) "Experience-based Discrimination: Classroom Games", *Journal of Economic Education* 36: 160-70.

Hogg, M.A. and Abrams, D. (2003) "Intergroup Behaviour and Social Identity", in *The Sage Handbook of Social Psychology*, ed. by M.A. Hogg and J. Cooper. London: Sage, pp. 407-423.

McGarty, C., Yzerbyt, V.Y. and Spears, R. (eds. 2002) *Stereotypes as Explanations: The Formation of Meaningful Beliefs about Social Groups*. Cambridge University Press.

Merton, R.K. (1957) "The Self-fulfilling Prophecy", in *Social Theory and Social Structure*. London: Free Press.

Richerson, P. and Boyd, R. (2001) "The Evolution of Subjective Commitment to Groups: A Tribal Instincts Hypothesis", in *The Evolution of Subjective Commitment*, ed. by R.M. Nesse. Russell Sage Foundation.

Rodgers, W.M. (ed. 2006) *Handbook of the Economics of Discrimination*. Cheltenham: Edward Elgar.

Schotter, A. and Weigelt, K. (1992) "Asymmetric tournaments, equal opportunity laws, and affirmative action: some experimental results", *Quarterly Journal of Economics* 107: 511-39

Sherif, M., Harvey, O., White, B., Hood, W. and Sherif, C. (1961) *Intergroup Conflict and Cooperation: The Robbers Cave Experiment*. Norman: University of Oklahoma Book Exchange.

Tajfel, H. (1982) *Human Groups and Social Categories*. Cambridge University Press.

Tajfel, H., Billig, M., Bundy, R. and Flament, C. (1971) "Social Categorization and Intergroup Behaviour", *European Journal of Social Psychology* 1: 149-178.

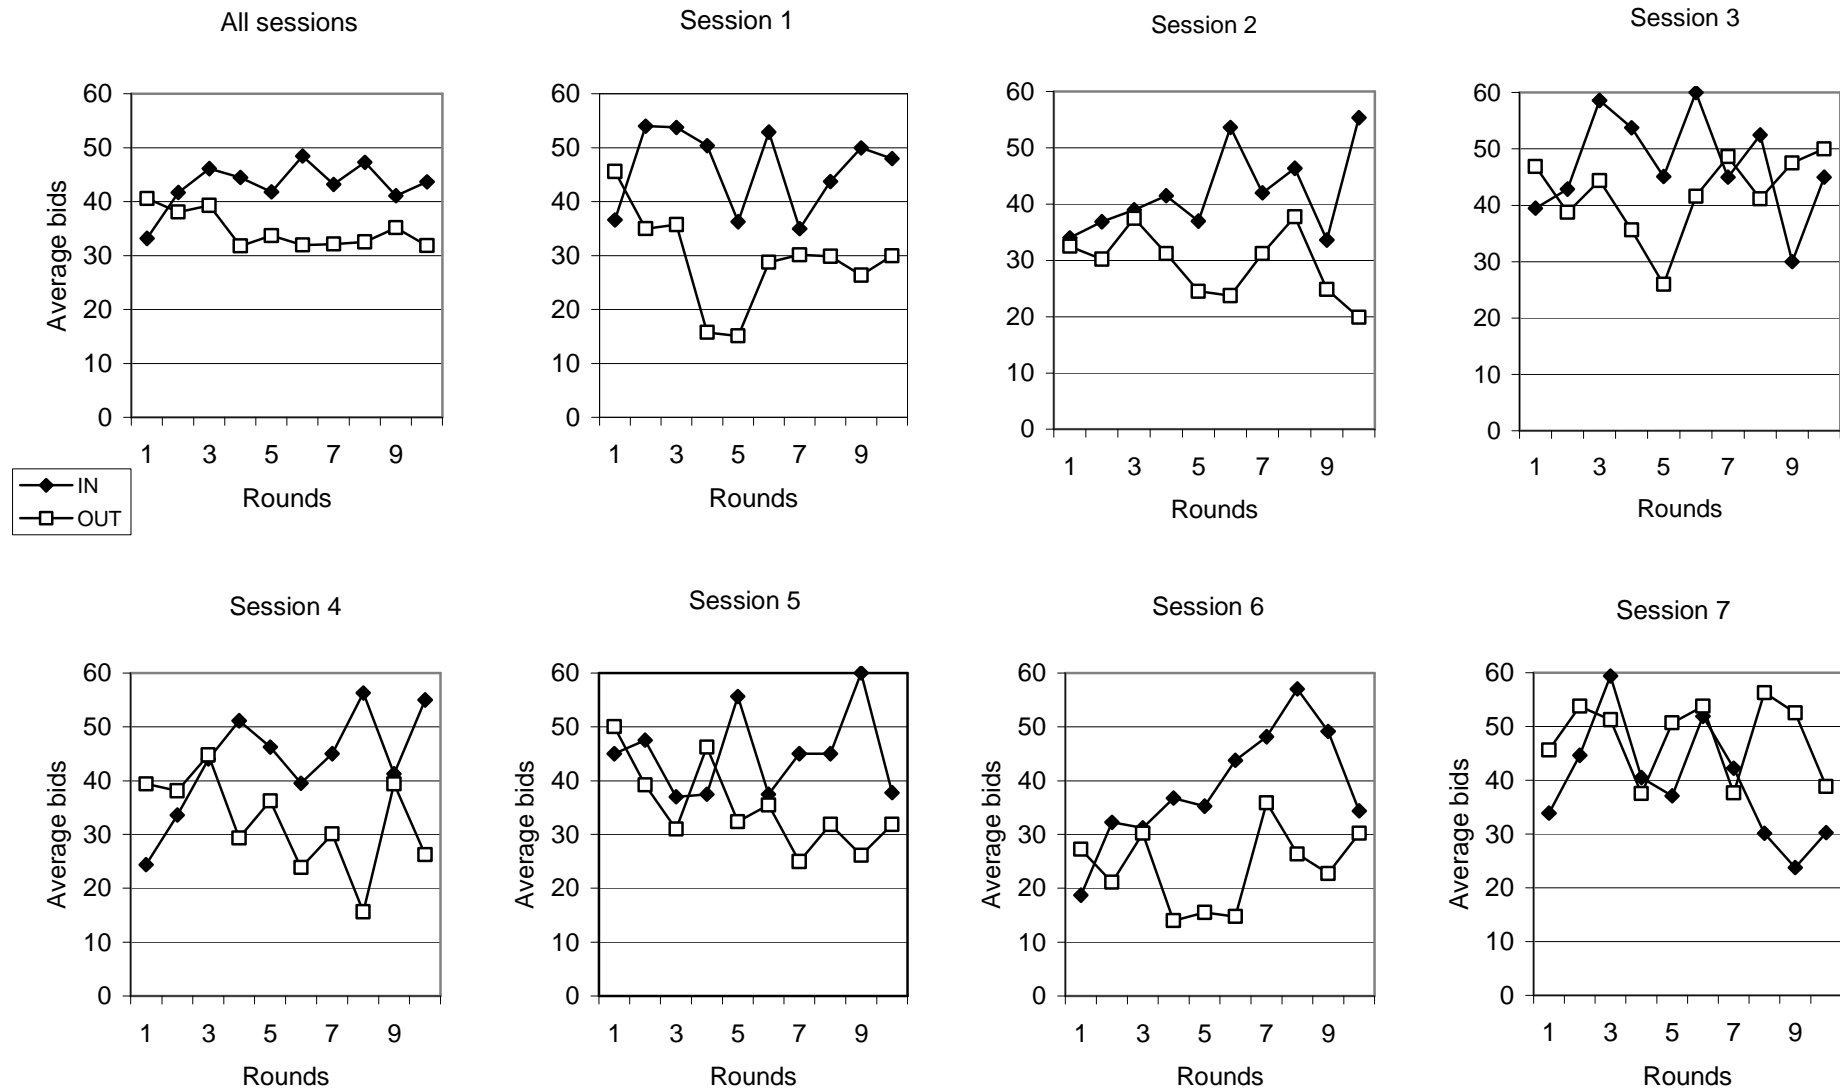


Figure 1: Average bidding patterns of in-group and out-group players, on aggregate (a) and per session (b-h).

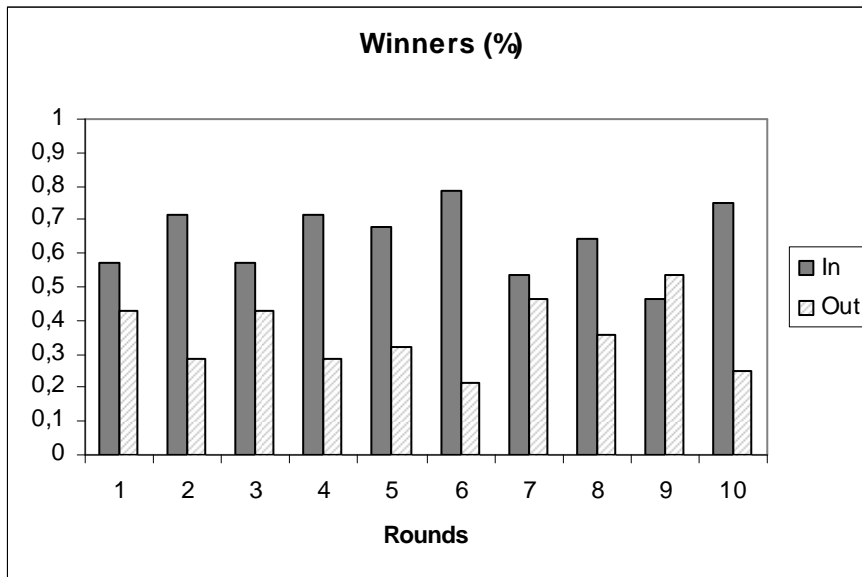


Figure 2: Proportion of winners per type of bidder, in each round.

FIGURE 3a, b, c

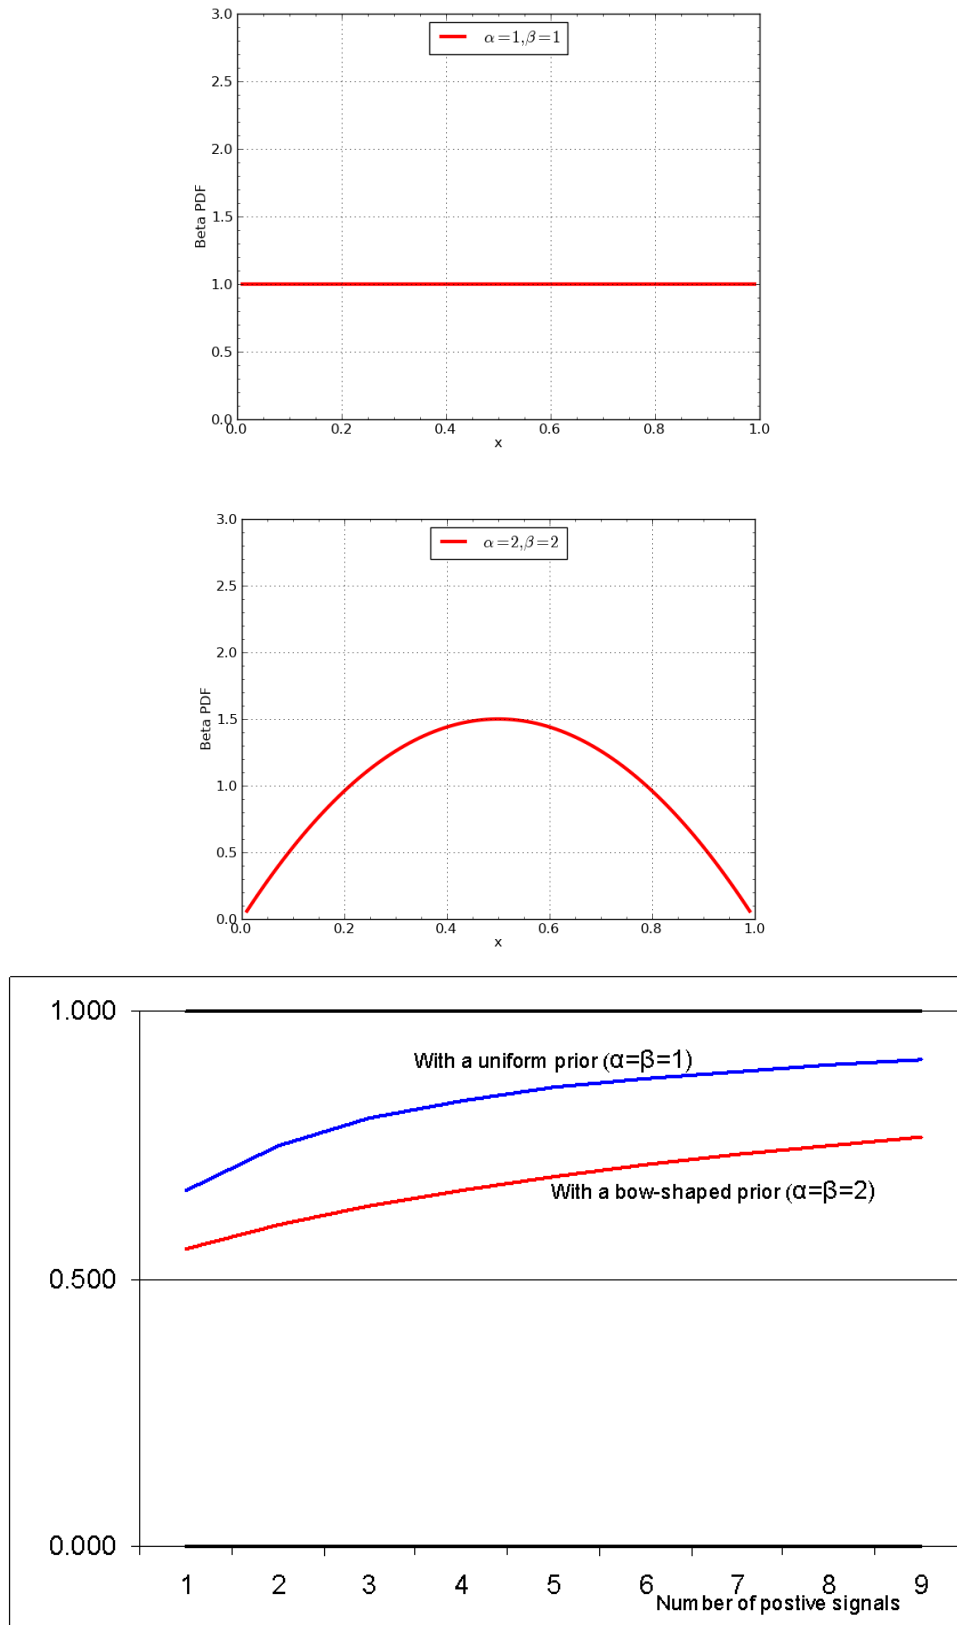


Figure 3: Two examples of prior belief distributions over auctioneers' decisions in tie-breaks (3a): uniform distribution with $\alpha = \beta = 1$, (3b): bell-shaped with $\alpha = \beta = 2$. The convergence properties of these distributions are represented in 3c: convergence is quicker with the uniform prior distribution.

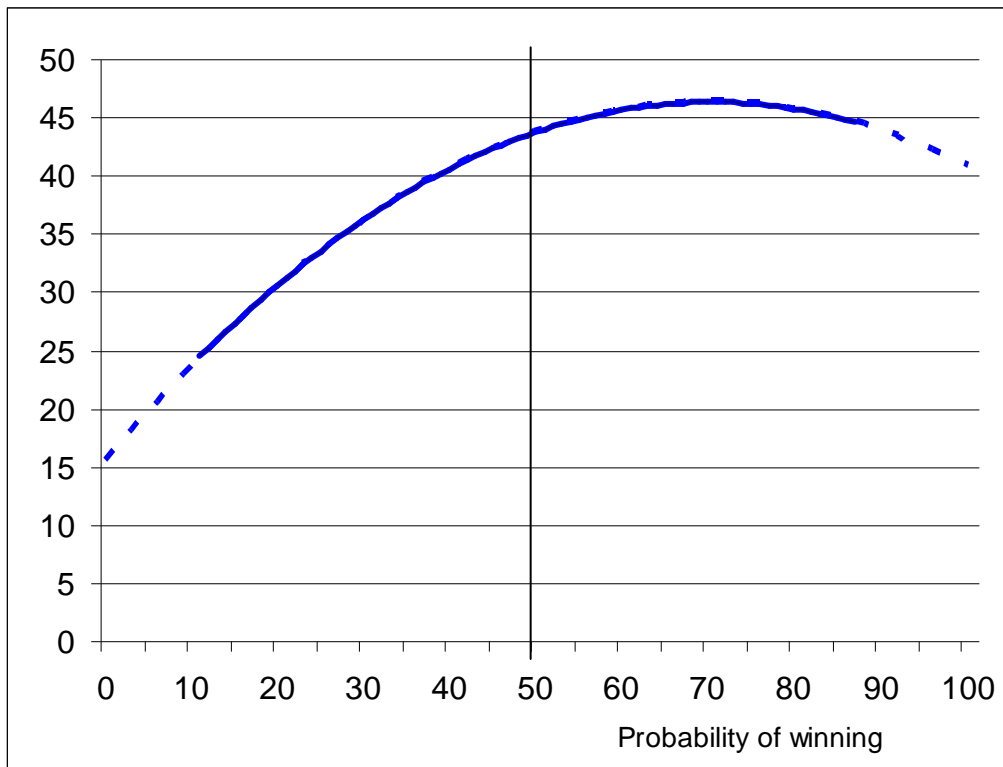


Figure 4: Bidders' reaction function, assuming a Beta(1,1) prior and posteriors that follow the stream of positive and negative meaningful signals. Winning likelihoods (expectations) are on the horizontal axis, the level of bids on the vertical axis. Dashed line for values of probability never observed in the sample.

	IN			OUT		
Round	0	0<bid<60	60	0	0<bid<60	60
1	0.05	0.70	0.25	0.02	0.59	0.39
2	0.09	0.46	0.45	0.16	0.44	0.41
3	0.10	0.30	0.60	0.20	0.34	0.46
4	0.12	0.33	0.55	0.20	0.42	0.38
5	0.14	0.27	0.59	0.32	0.34	0.34
6	0.11	0.14	0.75	0.34	0.20	0.46
7	0.18	0.18	0.64	0.28	0.27	0.45
8	0.14	0.15	0.71	0.32	0.22	0.46
9	0.25	0.18	0.57	0.34	0.18	0.48
10	0.13	0.23	0.64	0.36	0.18	0.46

Table 1: Polarization of bids in the course of the experiment

Highest bids made by		Number of winners		Total
IN	OUT	IN	OUT	
2	0	23	4	27
1	0	49*	5*	54
2	1	35	16*	51
2	2	15	4	19
1	1	41*	24*	65
1	2	13*	12	25
0	1	2*	27*	29
0	2	2	8	10
		64.3%	35.7%	280

Table 2: Highest bids and auctioneers' decisions, per type of player.

Dependent variable: Bid(<i>t</i>)	(1) unbiased prior	(2) biased prior
Bid(<i>t</i> -1)	.0175 (0.3)	.0288 (0.5)
Likelihood	.8728 (2.5)**	.8160 (2.3) **
Likelihood^2	-.0062 (1.9)*	-.0058 (1.8) *
Likelihood*out	-.3367 (1.5)	-.3219 (1.1)
Likelihood^2*out	.0037 (1.0)	.0049 (0.9)
Fixed effects	yes	yes
N	896	896
Test H ₀ =zero autocorrelation (Prob > z)	order 1: 0.0000*** order 2: 0.6795	order 1: 0.0000*** order 2: 0.6091

Notes: Absolute value t in parenthesis. * significant 10%; ** significant 5%

Table 3: Players' reaction function.

Appendix: Experimental instructions

You are about to take part in an experiment funded by several foundations for research purposes. During the experiment you will have the opportunity to make choices that will affect your earnings. The exact amount of your earnings will depend on your behavior and on the behaviour of other participants. All the money you earn will be paid when the experiment is finished. The experiment is divided in two stages. At the beginning of each stage you will receive instructions that will explain the task to be performed.

STAGE 1

In the first stage you will have to perform a collective task. At the beginning of the experiment each subject has been assigned a colored bracelet. Subjects with the red bracelet belong to the Red group, while those with blue bracelet belong to the Blue group.

In the first stage of the experiment each group will have the opportunity to earn 30 euro by solving a simple problem. If the group is successful the amount earned will be distributed among all group members equally, otherwise nobody will gain anything.

To earn the prize, the group must be able to memorize at least 10 verses of a poem and to write them correctly on a sheet of paper that will be provided by the experimenters. You can keep the text of the poem for five minutes, during which you will try to memorize a sufficient number of verses. At the end of this period, the assistant will collect the text of the poem, and will give you a blank sheet of paper and a pen that you can keep for another five minutes, during which you will write the memorized verses. At the end of this period the assistant will collect the paper and check the correspondence with the text of the poem.

Note that:

- You can report on paper more than ten lines (to maximize the probability of success).
- Only verses that are reported correctly will count (for example verses with a wrong word are not valid, while minor misspelling do not count).
- If less than ten verses are correct the gain of the group is zero.
- If the correct verses are eleven or more, the gain of the group is still 30 euro.

The amount earned by the group will be divided among all members in equal shares and paid at the end of the experiment, along with the money earned in the other phases of the experiment.

Are there any questions?

STAGE 2

From this moment the second stage of the experiment begins. The instructions we have distributed are only for your personal use. Unlike in the first stage, it is forbidden to communicate with other participants. If you have any doubts or questions, please ask the assistants. If you do not stick to this rule, we will have to exclude you from the experiment and from additional money rewards.

During the second stage of the experiment you will have the opportunity to make choices that will affect your earnings. The exact amount of your earnings will depend on your decisions and the decisions of other participants. The choices made by each subject, however, will be completely anonymous. The anonymity will be maintained both during and after the experiment: all the money you earn will be paid privately at the end of the experiment.

Description of the task

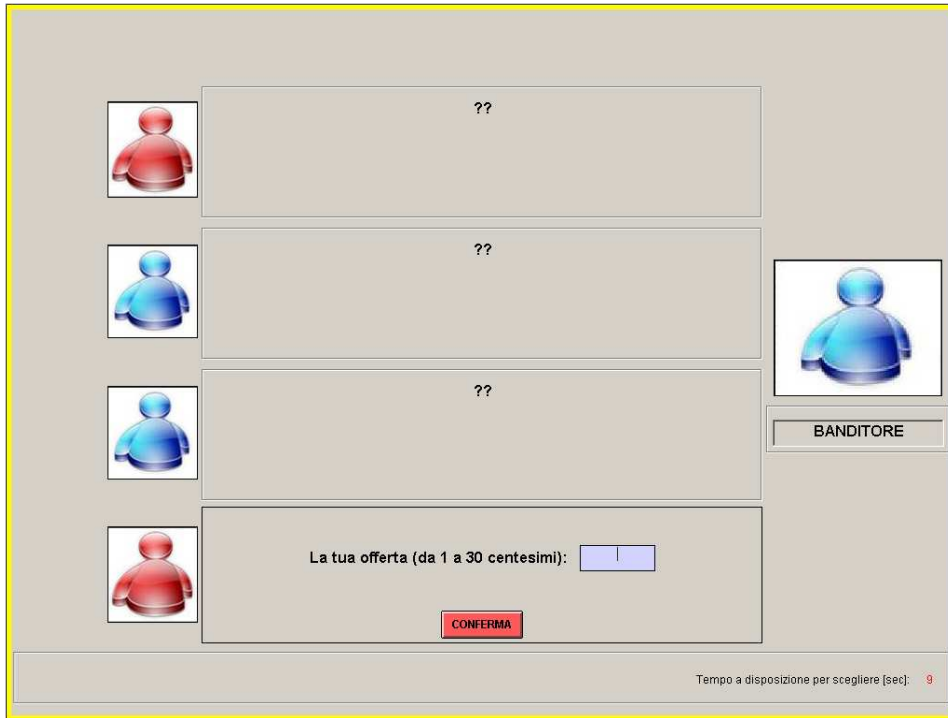
We introduce now the situation you will face at this stage. Twenty subjects will participate in each experimental session. During the experiment each subject will play with other subjects randomly selected among the participants, and identified only by a role (auctioneer or competitor) and a color (Red or Blue, depending on the bracelet that you wear) that represents the group affiliation of each experimental subject.

Each subject will participate in a series of ten sealed bid auctions. Four competitors (two Red and two Blue) and an auctioneer (Red or Blue) participate in every auction. In each of the ten rounds the players will be grouped randomly, so each time you will play with potentially different opponents. The latter will be identified only by their role and their color, while their personal identity will remain strictly confidential. The role of each player and her colour will remain fixed throughout this stage of the experiment.

At the beginning of each auction each competitor will have a budget of 60 cents to invest to win a prize worth 400 cents. The auctioneer will decide to assign the prize to one of the bidders. In order to win the prize, each of the four bidders may offer a sum of money that does not exceed her budget (up to 60 cents). Bids will be made by typing a number from 0 to 60 in a screen like the one found on the next page.

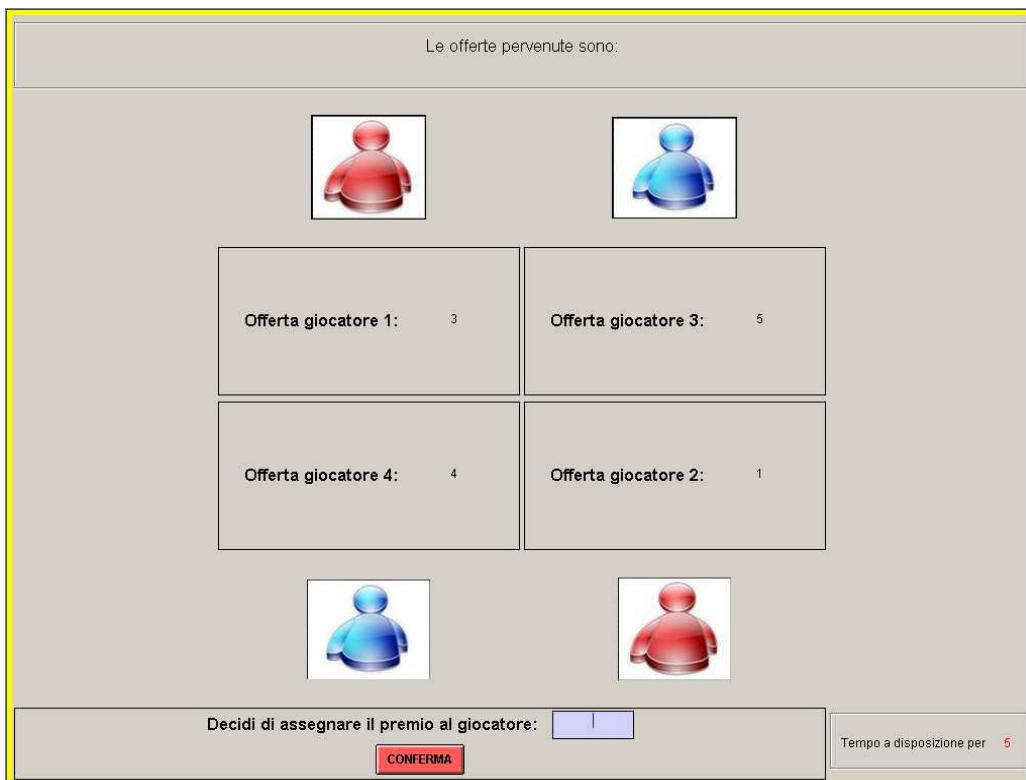
Bids will be made “in a sealed envelope”, i.e. simultaneously, without communicating with other players, and without the ability to operate subsequent corrections. Bids will not be reimbursed: the amount of money offered will be subtracted from the budget of any bidder regardless of whether she has won the auction or not. In other words:

- Gain of the winner = $60 - \text{bid} + 400$
- Gain of the other bidders = $60 - \text{offer} + 0$



The auctioneer may award the prize to any bidder. Each bidder will be identified on her screen only by a number and by the colour of her group (Red / Blue). (Note: Because the identity of competitors changes at every round, the subject identified as Player 1 in the first round will very likely not be the same person identified by the number 1 in the other rounds.)

The auctioneer will make his choice by typing the bidder's number in a screen like this:



The auctioneer will earn a sum of money equal to the offer of the bidder who was awarded the prize (i.e. the winning bid) multiplied by three. In other words:

- Gain of the auctioneer = (winning bid \times 3)

As mentioned, each subject will participate in ten subsequent auctions. At each round players will be grouped randomly, so that every subject might play with potentially different opponents in different rounds. At the end of each auction the colour of the winner will be communicated. Each participant will also be notified the amount of money earned in that auction.

At the end of the experiment we will ask you to fill a short questionnaire. Then the experimenters will pay privately the money that you have earned during the experiment. You will be asked to sign a receipt, and then to leave the room quietly. We would be grateful if you will not discuss the experiment with other students after leaving the laboratory.