Remembering faces and remembering whom to trust - An experimental study on cognitive racial discrimination
VERY PRELIMINARY AND INCOMPLETE

Michele Belot

October 21, 2011

Abstract

The question addressed in this study is whether limitations in the ability to remember people from other races can explain differences in trust decisions within and across race. I conduct an experiment in a controlled environment. Participants are presented with a set of pictures of people of different races - East Asian and Caucasian White - and each person is labelled as "trustworthy" or "not trustworthy" in a random manner. Subjects are then shown all pictures again, and are asked whether they trust each person or not. The question is whether subjects can better recall whom to trust and whom to distrust among people from their own race in comparison to other races. An asymmetry could provide an alternative explanation for why trust is easier to sustain within race than across race.

1 Introduction

Ethnically homogenous societies seem to enjoy higher levels of trust and social capital, as well as higher economic success[1]. The question is whether there is *causal* relationship running from ethnic diversity to trust and social capital and if that is the case, what is the mechanism? These questions are very difficult to address with observational data. Experimental studies offer a more promising route. Fershtman and Gneezy (2001) conducted an experimental study in

[1] See Knack and Keefer, 1997; Uslaner, 2000: 580; Alesina and La Ferrara, 2002; Paxton, 2002; Costa and Kahn, 2003; Helliwell, 2003; Hero, 2003

the Israel Jewish society. Participants play a trust game and learn about their partner's ethnicity implicitly through their name. The main question is whether trust decisions and trustworthiness are conditional on the ethnicity of the partners involved. The experiment was designed to disentangle between two mechanisms that could explain differences in trust behaviour: Preferences (Becker's (1961) taste-based discrimination) or stereotypes (Arrow's (1773) statistical discrimination). To do this, they compare behaviour in a trust game (that could be driven both by stereotypes and preferences) and behaviour in a dictator game (that can only be driven by preferences). Similar experiments were conducted by Bouckaert and Dhaene (2004) in Belgium with Turkish and Belgian small businessmen and Van Der Merwe and Burns (2008) in South Africa with Black and White particicipants.Interestingly, these studies generally fail to find evidence for an in-group bias in trust. Bouckaert and Dhaene find no differences in behaviour correlated with ethnicity. Fershtman and Gneezy and Van Der Merwe and Burns find a systematic mistrust with respect to one of the ethnic groups: One group is perceived as more trustworthy both by their own group and the other ethnic group, and this is mainly driven by behavioral stereotypes rather than preferences. These findings are interesting, but do not provide direct support for a causal relationship between ethnic heterogeneity and cooperation. As it is now, we do not know whether this causal relationship exists or not.

This paper aims at exploring an alternative mechanism for explaining differential levels of trust across ethnic groups. The starting point is that trust decisions often take place in a context of repeated interactions and often involve remembering people. Axelrod's (1971) seminal work on cooperation stresses that "the ability to recognize the other player from past interactions, and to remember the relevant features of those interactions, is necessary to sustain cooperation" (p. 139). In that context, bounded memory could play a role in explaining differences in cooperation within and across racial groups. The question we ask here is whether there is an own-race-bias in how we retrieve relevant economic information about people. Are we better at remembering who is trustworthy and who is not within our own race than across race?

It is well-known that people are better at remembering faces of people of their own race than other races (see Meissner & Brigham, 2001; Slone, Brigham, & Meissner, 2000 for reviews). Psychologists have repeatedly stressed the implications of such bias for criminal convictions based on eyewitness identification. Facial recognition plays a key role in many social interactions and therefore own-race biases in facial recognition may have implications in many social contexts. Often a face is all we rely upon to identify someone. Local merchants typically know their customers only by face, not by name or any other identifier. Professional conferences, business social events, interactions between teachers and students, or between students themselves, are all examples of environments with repeated encounters and where facial recognition is an essential technology used to identify others. This said, we obviously use other technologies to identify people - names for example.But all these technologies are likely to share common features, insofar as they are likely to be shaped by the environment we grow up in and by the nature of our social interactions.

In a trust context, the critical question is not whether one recalls a face or a name though. The relevant question is whether one recalls *who is trustworthy and who is not* - that is, whether one can retrieve *information* associated with a face or a name - and whether there are biases in the accuracy of recall of such information. Hanley (2008) documents the many situations in which people report being able to remember faces but cannot recall the context. Thus, we do not know whether better facial recognition implies better retrieval of information associated with people and more efficient decisions.

The idea that an own-race-bias in memory may affect cooperation and trust has been unexplored in economics. The implications are similar in nature to those derived in contexts of asymmetric information and noisy signals. A bias in memory implies that beliefs about trustworthiness are noisier across race than within race. Since cooperation is more difficult to sustain in noisier environments, this mechanism could potentially explain differences in cooperation rates and trust within and across races. Fryer and Jackson (2008) propose a model

of discrimination precisely based on bounded memory.[2] In their model people are sorted into categories and each category has a prototype - a unique vector of attributes. People keep track of the variation in attributes *across* categories but not within a category. People sorted into the same category are blended together. They argue that minority groups may be sorted into coarser categories than majority members because they are less likely to be involved in frequently repeated interactions. The implication is that the stored payoff-relevant information is less precise for minority groups, which results in a compression in rewards. In our context, the implications are similar. If we cannot keep track of who is trustworthy and who is not, we should be less likely to trust those who are trustworthy but from another race than our own. But this also means that we should also be more likely to avoid non-trustworthy people from our own race than the other race. This is precisely why, in a repeated game context, this mechanism may lead to higher levels of trust within race than across.

We conduct a simple experimental study aimed at answering a simple question: Are we better at keeping track of who is trustworthy and who is not within our own race? We study this question with a novel design, which involves a simple individual decision in a controlled repeated environment. In a first stage, we attribute at random the labels "trustworthy" and "not trustworthy" to people who are not involved in the experiment. Participants see pictures of these people, together with their trustworthiness label. In a second stage, participants have to decide whether to engage in a transaction with each person. They see the pictures of people again, but without their label, and are asked to choose for each person whether to (1) trust the person, (2) avoid the person, (3) play a lottery. The trust / avoidance decisions provide a fixed benefit of £0.75 if the person is indeed trustworthy / not trustworthy and £0 otherwise; and the lottery provides the same fixed benefit of £0.75 with a given probability p (which we vary across treatments) and £0 with probability $(1-p)$. We introduce the lottery option in addition to the "trust" decisions precisely to be

[2] This work fits in the literature on bounded memory, and coarse thinking see Mullainathan (2002) and Mullainathan et al. (2008).

able to capture levels of uncertainty in inferred trustworthiness. We introduce 2 treatments, with $p = \{0.6; 0.75\}$. In both cases, it is always better to choose the lottery if one does not remember anything. Better memory on the other hand should increase the proportion of people "trusted" and "avoided". The pictures shown in the second stage include some new faces of people participants have not seen before. We do this to be able to disentangle between different sources of confusion: A first source of confusion is the own-race-bias in face recognition identified by psychologists - two faces are confused with each other. The second source is the inability to map a face to payoff-relevant information, even if the face is correctly identified. This second source of confusion cannot be present for new faces, while both sources of confusion are possible for faces that were shown in the first stage.

We chose a non-strategic decision design because it allows us to isolate the role of memory from the alternative mechanisms we discussed earlier. First, we can strictly limit the consequences of decisions to the decision-maker: the decision only affects the participant's own payoff and does not affect the person who is the object of the decision. The people appearing on the pictures do not incur any loss or gain by being selected or not. This rules out a role for preferences and other-regarding considerations, such as fairness or willingness to provide benefits to own-group members over others. Second, the payoffs associated with each candidate are fully controlled for and participants are fully informed about these individual payoffs and about the procedure of assignment of labels to candidates. This exogenous assignment means that the faces of candidates do not contain any relevant information and leave little room for stereotypes. This is important given the recent evidence showing relationships between facial features (and in particular race) and inferences about personal characteristics such as competence and trustworthiness (Todorov et al. (2005), Eckel and Petrie (2008), Rule and Ambadie (2008), Todorov and Duchaine (2008), Duarte et al. (2009)). Here the only cognitive mechanism that can correctly map faces to values is memory.

The rest of the paper is structured as follows. Section 2 discusses the relevant

literature in psychology. Section 3 presents the experimental design and Section 4 the analysis and results. Finally, I conclude in Section 5, by discussing relevant applications and new research questions.

2 The Own-Race Bias in Facial Recognition

There is a well-established literature in psychology on the "own-race bias" in facial recognition - the fact that people are better at remembering faces of people of their own race than other races (see Meissner & Brigham, 2001; Slone, Brigham, & Meissner, 2000 for reviews). Psychologists have repeatedly stressed the implications of such bias for criminal convictions based on eyewitness identification. However, as we mentioned earlier, in a social or economic context, the critical question is not whether one recalls a face or a name. The relevant question is whether one recalls *who a person is* - that is, whether one can retrieve relevant *information* associated with a face or a name - and whether there are biases in the accuracy of recall of such information. Hanley (2008) documents the many situations in which people report being able to remember faces but cannot recall the context. We do not know how people memorise payoff-relevant information and map this information to the corresponding identifiers. As we shall argue later, it is not clear at all that own-race biases should be as large or even arise. But obviously, the findings on the own-race bias in facial recognition are relevant for our study, so we review this literature in some detail.

Psychologists and neuropsychologists have extensively studied the cognitive and neurological processes involved in facial recognition (see Duchaine (2008) for a review). The own-race bias (ORB) is one of the most robust empirical findings in the literature on facial recognition. Meissner and Brigham (2001) provide a detailed meta-study of the last thirty years of literature, reviewing 39 articles involving the responses of over 5,000 participants. The overwhelming consensus among social psychologists is that an own-race bias exists and is quite large.

The cognitive and social factors responsible for the ORB remain unclear

(Slone et al., 2000). Theories proposing that the degree of interracial contact should be negatively associated with level of ORB have been only weakly supported (Chiroro & Valentine, 1996). Meissner & Brigham (2001) show in their meta-analysis that interracial contact accounts for only about 2% of the variance in ORB across samples. Although negative racial attitudes are correlated with limited interracial contact, no relationship has been found between the ORB and racial attitudes, whether explicit or implicit (Ferguson, Rhodes, & Lee, 2001). Training does not seem to help much either. Lavrakas et al. (1976) show that training could reduce the magnitude of the ORB, but the effect was short-lived: One week later there was no difference between trained and untrained participants.

Some evidence suggests that one reason for the ORB may be that cross-race faces are processed differently than own-race faces. In essence, cross-race faces may be perceived more "holistically" - more like objects (Rhodes et al., 1989; Tanaka et al., 2004). This idea is confirmed by neurophysiological studies. Neurophysiologists have identified specific areas of the brain active in the processing of faces and that the processing of cross-race faces is different from the processing of own-race faces (Golby et al. (2001), Cunningham et al. (2004) and Duchaine (2008) for a recent survey).

Levin (1996, 2000) proposes that other race effects are caused by selection of different facial features in same and other race faces. Whereas individuating information is selected in same race faces, race specifying information is emphasized in representations of other race faces at the expense of individuating information. In fact, race has been shown to be one of the prime characteristics encoded in human interactions, together with gender and age (Montepare and Opeyo (2002)). On the other hand, a number of studies show that faces rated as distinctive are more accurately remembered (Shepherd et al. (1991) and Valentine, 1992 for a survey of relevant studies). Thus, in a situation where race is a scarce attribute, it could serve as an obvious marker of identity and improve recognition significantly.

Recent work suggests that face recognition develops with age. Pascalis O.

et al. (2002) showed that 6-month old infants, 9-month old infants and adults were able to discriminate between human faces but only 6-month-olds could discriminate between monkey faces. This phenomenon is similar to the loss of sensitivity to phonemes not used in the infant's native language (Werker and Tees (1984), Kuhl (1992), Aislin et al (1998)). Differential processing of faces of different races follows a similar developmental course.

Note that these findings are not directly relevant to the question we are interested in here. To understand the economic implications of cognitive biases in re-identification, one needs to study the joint recall of identities and payoff-relevant information - how people record and retrieve information of the type "Person x has productivity y". The most relevant studies in psychology are those that study the recall of *associations* between faces and information. The seminal work in that area is Taylor et al. (1978), who study how participants recall the contents of interactions between people of mixed gender and race ("*Who said What?*")[3]. They show that participants are more likely to misattribute statements of people of the same race than different races. To the best of our knowledge, there are no studies providing evidence on how people memorize identities and payoff-relevant information. Payoff-relevance introduces a specific nature to information and perhaps commands a specific way of recording and organising information.

3 Experimental Design

The experiment was conducted at the laboratory of the Nuffield Centre for Experimental Social Sciences (Oxford) in October 2011. @@ participants were recruited: @@ Caucasians and @@ East Asians. Invitations were sent by e-mail to participants in the pool with East Asian and British last names, without mentioning race or ethnicity[4]. Sessions were also relatively small in size (maximum 15 participants at a time) and the split across race was not equal for

[3] I am very grateful to Oliver Curry for informing me about this literature

[4] Participants were recruited using ORSEE (Greiner, 2004). The invitation asked for participants between the ages of 18 and 30. We have excluded participants above 50 years old from the analysis (n = 3).

each session. The experiment lasted for 30 minutes in total, including a post-experimental questionnaire asking information about ethnicity, age, occupation, country of birth, age of arrival in the UK and a self-assessment of ability to remember faces of people in general. Participants received a £4 show-up fee and an additional payment depending on the performance in the memory task (see explanation below), bringing the total payment to @@ on average.

The experiment is structured in 2 stages. The first stage is a *viewing stage* where participants view an automated sequence of 16 pictures. Each picture is shown for 5 seconds and then the screen moves to the next picture.[5] Each picture appears together with its label. We assign labels to each picture in a random manner, stratified by gender and race. There are 4 pictures for each gender and race combination. We always ensure that 50% in each group are assigned the label "trustworthy" and the other 50% the label "untrustworthy". .Participants are not informed about the stratification (we do not mention race or gender at any point), but are informed that 50% of the population is labelled as trustworthy.

The second stage is a *selection stage* – participants see the pictures of 16 people again, but without their associated label, and are asked for each candidate to choose between 3 options:(1) trust this person, (2) distrust this person or (3) play a lottery. Importantly, pictures appear in a different sequence, and from a different angle. We introduce 4 pictures of new people, among the 16, who were not shown in stage 1. These new people are also implictly attributed a label "trustworthy" or "not trustworthy" with a probability 0.50. The sequence is not automated in the second stage and participants can go back and forth between pictures for 3 minutes. Note that the sequence of presentation of candidates is randomized for each participant, and for each stage.

Participants earn £0.75 if they correctly trust (avoid) people who are trustworthy (not trustworthy), including for new people. If they choose to play the lottery, they receive £0.75 with probability p, $p = \{0,6; 075\}$. Their total earn-

[5] The choice of number of pictures, time and mix of gender is in line with the common practice in psychology studies.

ings are equal to the sum of the individual earnings for all 16 pictures shown in the second stage.

The instructions have been written carefully to inform participants in detail about the procedure of assignment of labels to candidates (see Appendix). The instructions do not mention race at any point.

Pictures of candidates

Pictures of candidates were drawn from a database provided by TARRLAB [6] These pictures show only the face of the person. Pictures were selected according to a number of criteria to guarantee homogeneity in shooting conditions[7]. The database contains 11 East Asian men, 15 East Asian women, 35 Caucasian men and 44 Caucasian women.

For each subject, a set of 24 candidates was randomly chosen. Two pictures of each candidate from a different angle were randomly chosen (one used in the viewing stage and the other used in the selection stage). This ensures that the task involves face recognition rather than picture recognition and prevents participants from using other cues than the face itself to remember the person. The sequence of viewing is determined randomly for each participant, for both the viewing and selection stages.

A picture of a mixed race person has been chosen to illustrate the instructions (see instructions in the appendix).

4 Analysis

4.1 Predictions

It is useful to outline a number of implications of different models that could drive selection decisions in this context. I describe the implications associated

[6] Face-Place Face Database Project (http://www.face-place.org/); Copyright 2008, Michael J. Tarr. Funding provided by NSF award 0339122.

[7] All images were extracted from standard digital video (720x480), with the background removed and the faces scaled to be roughly equated in terms of size. The pictures were selected according to the following parameters: Race (East-Asian and Caucasian White); shave/stuble no make-up; no beards or mustache; no facial hair or visible make-up; no glasses; natural hair (no wig); neutral affect; orientations: 0°, 15° left, 15° right, 30° left, 30° right.

with a memory-based model in the spirit of Fryer and Jackson (2008) and I contrast these implications with those from models based on preferences or stereotypes. I am specifically interested in implications regarding the probability of entering the selection and how this probability varies with the value attached to the picture.

Bounded memory

The simplest way to describe the decision problem is as follows. Participants see a sequence of candidates with labels $x_i = \{trustworthy, untrustworthy\}$.They know that 50% of the population is trustworthy. Let us assume that the decision maker forms categories of people based on observable facial attributes: for example, according to race, gender, colour of the hair or eyes. Suppose this results in K equal size categories for the own race and J equal size categories for the other race, with $J < K$. The decision maker cannot distinguish between people sorted into the same category.

In addition to keeping track of the observable attributes uniquely identifying each category, the decision maker also needs to keep track of the payoff-relevant information. Let us assume that she keeps track of the proportion of people who are trustworthy in each category.

Each category is then associated with an expected value $\widehat{q}$, which corresponds to the proportion of trustworthy people in the category. In this framework, the optimal decision (under risk neutrality) consists in:

- trusting all people in the category if $\widehat{q} > p$
- distrusting all the people in the category if $\widehat{q} < (1-p)$.
- play the lottery if $(1-p) > \widehat{q} > p$

For example, suppose p is 60%. Then the optimal decision is to trust if $\widehat{q} > 60\%$, to distrust if $\widehat{q} < 40\%$ and to play the lottery if $40\% > \widehat{q} > 60\%$.

If memory biases lead to less precise beliefs, we would expect trust decisions to be more uncertain across race than within race and people to choose the lottery more often when it involves a person from the other race than their own.

It is important to point though that bounded memory does not necessarily

lead to biases in initial trust decisions. Suppose for example that categories are finer for one's own race such that there are 2 categories of 3 people above 50%, one with $\widehat{q} = 70\%$ and the other with $\widehat{q} = 55\%$ but only one category of the other race with $\widehat{q} = 65\%$. Then the optimal decision under risk neutrality is to trust all people from the other race belonging to the category above 50%, but trust only those from the first top category within one's own race. Thus, own-race biases in memory do not necessarily lead to in-group biases. But one clear implication is that the proportion of inefficient decisions - on the one hand, trust and lottery decisions involving non-trustworthy people and, on the other hand, avoidance and lottery decisions involving trustworthy people - should be larger across race than within race. And this is precisely this mechanism that could explain why cooperation rates differ in a repated game environment.

Preferences or Stereotypes

The experiment is designed to isolate the role of memory from other possible mechanisms that could link trust and ethnicity. First, preferences and tastes should have a limited role here by design since there is no "interaction" between the decision-maker and the person who is involved in the decision (the candidate on the picture). There is no other benefit from selecting a particular picture other than than the economic value attached to that picture. Also, decisions only affect the decision-maker and not the candidates, such that other-regarding considerations should not affect decisions. Yet participants come to the lab with a real world experience. It could be for example that they associate people from other races with untrustworthiness. The implication would be that people from other races are less likely to be trusted. So we should find a lower proportion of people "trusted" and a larger proportion of people "avoided", which is a different prediction than with bounded memory. Second, the design leaves little room for stereotypes (Phelps, 1972 and Arrow, 1973) to play a role. Stereotypes can affect decisions in environments where there is uncertainty about the value of a person. Here the environment rules out this uncertainty by design: Participants are perfectly informed about the trustworthiness of the people on the pictures,

Table 1: Treatments and number of participants

Caucasian White participants	
East-Asian participants	

Table 2 - participants summary statistics

	Caucasian White participants	East-Asian participants
Age		
Share women		
Mean earning		

and about the distribution of trustworthiness in the population. Yet, again, participants may bring stereotypes formed outside the lab into the lab. The implications of negative stereotypes are identical to the implications of same race preferences. They also come down to attributing a lower level of trustworthiness to people of the other race and should also lead to a lower probability of being trusted and a higher probability of being "avoided". Again, these implications contrast with the implications of a memory-based model.

4.2 Analysis and results

[The results will be added by Tuesday 25 October - I outline here how I intend to present the results]

4.2.1 Summary statistics

Tables 1 and 2 present summary statistics of the participants.

4.3 Analysis

* Proportion "trust", "avoid" and "play the lottery" by participant ethnicity and picture ethnicity (by treatment and depending on whether it was a new face or not)

* Proportion of mistakes: For faces shown in the first stage: proportion of people trusted while they are not trustworthy, avoided while they are trustworthy. For new faces: Proportion trusted/avoided

5 Discussion and conclusion

[To be added]

References

[1] Aigner, D.J. and G.G. Cain (1977), Statistical theories of discrimination in labor markets, *Industrial and Labor Relations Review* 30 (2), 175-187.

[2] Aslin, R. N., Jusczyk, P. W., & Pisoni, D. B. (1998). Speech and auditory processing during infancy: Constraints on and precursors to language. In D. Kuhn & R. Siegler (Eds.), *Handbook of child psychology: Cognition, perception, and language* 2, 147–254, New York: Wiley.

[3] Alesina, A. and La Ferrara, E. (2002). Who trusts others? Journal of Public Economics, 85, 207–234.

[4] Axelrod, R. and W.D. Hamilton (1981), The Evolution of cooperation, *Science* 211, 1390–96.

[5] Arrow, K. (1973), The Theory of discrimination, in *Discrimination in Labor Markets*, O. Ashenfelter and A. Rees, eds. Princeton, NJ: Princeton University Press.

[6] Basu, S., J. Dickhaut, G. Hecht, K. Towry and G. Waymire (2009) Record-keeping alters economic history by promoting reciprocity, *Proceedings of the National Academy of Science* 106(4), 1009-14.

[7] Becker, G. S. (1961), The economics of discrimination, 2nd Ed. Chicago: University of Chicago Press.

[8] Van Der Merwe, W.G. and J. Burns (2008), What's in a name? Racial Identity and Altruism in Post-Apartheid South Africa, South African Journal of Economics 76 (2), 266–275

[9] Chiroro J and T.Valentine (1996), An investigation of the contact hypothesis of the own-race bias in face recognition. *Quarterly Journal of Experimental Psychology.*48A, 879–894.

[10] Cornell, B. and I. Welch (1996), Culture, information, and screening discrimination, *The Journal of Political Economy* 104(3) 542-571

[11] Costa, D. L and Kahn, M. E. (2003). Civic engagement and community heterogeneity: an economist's perspective. *Perspectives on Politics* 1, 103–111.

[12] Currarini, S., M.O. Jackson & P. Pin (2009), An economic model of friendship: Homophily, minorities, and segregation, *Econometrica*, 77(4), 1003-1045.

[13] Cunningham W.A., M.K. Johnson, C.L. Raye. J.C. Gatenby, J.C. Gore and M.R. Banaji (2004), Separable neural components in the processing of Black and white faces. *Psychological Science.* 15, 806–813.

[14] Duarte, J., S. Siegel and L.A. Young (2009), Trust and Credit. Available at SSRN: http://ssrn.com/abstract=1343275

[15] Duchaine,B., Yovel,G. (2008). Face Recognition. in Basbaum,A., R. Hoy, A. Kaneko, G. Shepherd and G. Westheimer (ed.) *The Senses: A Comprehensive Reference.* Amsterdam: Elsevier, 1st edition, 329-358. ISBN: 0-12-639482-2.

[16] Eckel, C. and R. Petrie (2010), Face Value, *American Economic Review* forthcoming.

[17] Ferguson D.P., G. Rhodes and K. Lee (2001) They all look alike to me: Prejudice and cross-race face recognition. *British Journal of Psychology*;92:567–577.

[18] Fershtman, C., and Gneezy, U. (2001) Discrimination in a Segmented Society: An Experimental Approach, *Quarterly Journal of Economic*s 116 (1), 351-377.

[19] Fong E and W.W. Isajiw (2000), Determinants of friendship choices in multiethnic society, *Sociological Forum* 15, 249–269.

[20] Fryer, R. and M.O. Jackson (2008), A Categorical Model of Cognition and Biased Decision Making, *The B.E. Journal of Theoretical Economics* (Contributions) 8 (1), Article 6.

[21] Gambetta, D. and M. Bacharach (2001), Trust in signs, in K. Cook (ed.) *Trust and Society*, New York: Russell Sage Foundation, 148-184.

[22] Golby A.J., J.D. Gabrieli, J.Y. Chiao and J.L. Eberhardt (2001). Differential responses in the fusiform region to same-race and other-race faces. Nature Neuroscience 4, 845–850.

[23] Greiner, B. (2004), The Online Recruitment System ORSEE 2.0 - A Guide for the Organization of Experiments in Economics. University of Cologne, *Working Paper Series in Economics* 10.

[24] Hanley, J.R. and G. Cohen (2008), Memory for people: Faces, names and voices in Cohen, Gillian and Conway, Martin A., *Memory in the real world*, Psychology Press, UK

[25] Ito T.A., G.R. Urland (2003), Race and gender on the brain: Electrocortical measures of attention to the race and gender of multiply categorizable individuals. *Journal of Personality and Social Psychology* 4, 616–626.

[26] Knack, S. and P. Keefer (1997), Does Social Capital Have an Economic Payoff? A Cross-Country Investigatio, *The Quarterly Journal of Economics* 112 (4): 1251-1288.

[27] Kuhl, P.K. (1993), Early linguistic experience and phonetic perception: Implications for theories of developmental speech perception. *Journal of Phonetics* 21, 125–139.

[28] Kurzban, R., J. Tooby and L. Cosmides (2001). Can race be erased?: Coalitional computation and social categorization. *Proceedings of the National Academy of Sciences*, 98(26), 15387-15392Lavr

[29] Lavakas P.J., J.R. Buri and M.S. Mayzner (1976), A perspective on the recognition of other race faces. *Perception & Psychophysics* 20, 475–481.

[30] Levin D.T. (1996), Classifying faces by race: The structure of face categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 22, 1364–1382.

[31] Levin D.T.. (2000) Race as a visual feature: Using visual search and perceptual discrimination tasks to understand face categories and the cross-race recognition deficit. *Journal of Experimental Psychology*: General 129, 559–574.

[32] Maclin O.H. and R.S. Malpass (2003), The ambiguous-race face illusion. *Perception* 32, 249–252.

[33] McPherson M, L. Smith-Lovin and J.M. Cook (2001). Birds of a feather: homophily in social networks. *Annual Review of Sociology* 27, 415–444.

[34] Meissner C and J. Brigham (2001), Thirty years of investigating the own-race bias in memory for faces. *Psychology, Public Policy, and Law.*7, 3–35.

[35] Mollica, k.A., B. Gray and L.K. Treviño (2003), Racial homophily and its persistence in newcomers'social Networks, *Organization Science*14 (2), 123-136.

[36] Montepare JM, Opeyo A (2002). The relative salience of physiognomic cues in differentiating faces: A methodological tool. *Journal of Nonverbal Behavior*. 26, 43–59.

[37] Mullainathan, S. (2002), A memory-based model of bounded rationality, *The Quarterly Journal of Economics* 117(3), 735-774.

[38] Mullainathan, S., J. Schwartzstein and A. Shleifer (2008), Coarse thinking and persuasion, *The Quarterly Journal of Economics* 123(2), 577-619, 05.

[39] Ng W, Lindsay R. (1994) Cross-face recognition: Failure of the contact hypothesis. *Journal of Cross-Cultural Psychology* 25, 217–232.

[40] Pascalis, O., M de Haan, C.A. Nelson (2002), Is face processing species-specific during the first year of life? *Science*.296 (5571), 1321 - 1323

[41] Paxton, P. (2002). Social capital and democracy: an inter-dependent relationship. *American Sociological Review*, 67, 254–277

[42] Phelps, E. (1972), The statistical theory of racism and sexism, *The American Economic Review* 62(4), 659-661

[43] Rhodes G, S. Brake, S. Tan and K. Taylor (1989) Expertise and configural coding in face recognition. *British Journal of Psychology* 80, 313–331.

[44] Rule, N. O. and N. Ambady (2008). The face of success: Inferences from chief executive officers' appearance predict company profits. *Psychological Science*, 19, 109-111.

[45] Shepherd, J.W., F. Gibling and H.D. Ellis (1991), The Effects of distinctiveness, presentation time and delay on face recognition, *European Journal of Cognitive Psychology 3,* 137-45.

[46] Slone A, J. Brigham and C. Meissner (2000), Social and cognitive factors affecting the own-race bias in Whites. *Basic and Applied Social Psychology*;22:71–84.

[47] Tanaka J, M. Kiefer and C.M. Bukach (2004), A holistic account of the own-race effect in face recognition: Evidence from a cross-cultural study. *Cognition* 93, B1–B9.

[48] Taylor, S., S. Fiske, N. Etcoff, and A. Ruderman (1978) Categorization and contextual bases of person memory and stereotyping. *Journal of Personality and Social Psychology* 36, 778-793.

[49] Tibbetts, E. A. and J. Dale (2007), Individual recognition: It is good to be gifferent, *Trends in Ecology and Evolution* 22 (10), 529-37

[50] Todorov, A., A.N. Mandisodza, A. Goren and C.C. Hall (2005), Inferences of competence from faces predict election outcomes. *Science*, 308, 1623–1626.

[51] Todorov, A. and B. Duchaine. (2008). Reading trustworthiness in faces without recognizing faces. *Cognitive Neuropsychology* 25, 395-410.

[52] Uslaner, E. M. (2000). Producing and consuming trust. *Political Science Quarterly*, 115, 569–590

[53] Valentine, T. (1992), Towards an exemplar model of face processing The effects of race and distinctiveness, *The Quarterly journal of experimental psychology. A, Human experimental psychology* 44 (4) 671-703.

[54] Van Der Merwe, W.G. and J. Burns (2008), What's in a name? Racial Identity and Altruism in Post-Apartheid South Africa, South African Journal of Economics 76 (2), 266–275

[55] Werker, J.F. and R.C. Tees (1984), Phonemic and phonetic factors in adult cross-language speech perception. *Journal of the Acoustical Society of America* 75, 1866–1878

[56] Zak, P. and S. Knack (2001), Trust and growth, .111 470 *The Economic Journal*, , 295–321