GROUP IDENTITY IN INTERMEDIATED INTERACTIONS: LESSONS FROM A TRUST GAME WITH DELEGATION IN SOUTH AFRICA

Hanjo Hamann and Nicky Nicholls

ABSTRACT

We investigate the role of group identity in delegated decision-making. Our framework considers the impact of group identity (based on racial segregation in post-Apartheid South Africa) on decisions to appoint a representative in a trust game with delegated decision-making, where information on the race group of other players is either common or private knowledge. We test our framework experimentally on a sample of young South Africans who had never been exposed to experimental economics research. By exogenously matching parties according to their race group, we observe their endogenous trust and delegation behavior. Our results suggest that white players try to use information about group identity to increase profits, albeit unsuccessfully. This may help to explain distrust and coordination failures observed in real-life interactions.

Keywords: Experiment; agency; delegation; trust; group identity; race; South Africa

JEL classifications: C91; D02; D03; D83; J15; Z13

Experimental Economics and Culture
Research in Experimental Economics, Volume 20, 227–264

ISSN: 0193-2306/doi:10.1108/S0193-230620180000020008

1. INTRODUCTION

This chapter adds to the literature on delegated decision-making in situations in which trust and reciprocity are important. In particular, we ask how group identity, which has been widely noted to impact trust and reciprocity in nondelegated decisions, would impact delegation decisions. Consider some relevant delegation settings, some of which are widely noted in studies of delegated decision-making: (1) Division heads in firms are assigned a budget with which they have significant autonomy in making compensation decisions (reciprocity may translate into increased effort from workers) on behalf of the firm owner/shareholders. (2) Owners of rental properties appoint rental estate agents who have autonomy in determining appropriate rental rates and in selecting tenants (in anticipation of care being taken of the property and payments being made as agreed). (3) Sales representatives decide which potential buyers can be trusted with test-driving a car (in anticipation of the car not being stolen and of potential profit from selling the car).

Other settings we note where the group identity of the representative might be particularly relevant include policing in situations when race might appear as a salient group identifier: Consider by way of example the case of Ferguson in the USA, where a white police officer fatally shot a black teenager in 2014. Here, the (predominantly white) city council needed to appoint a head of police to represent them in engaging with the (predominantly black) population. Once appointed, the head of police (who would either mirror the racial identity of the city council or that of the population) would have significant autonomy in deciding how to interact with the population, who might reciprocate through increased positive engagement with police and city council. As another example, peacekeeping forces sent to help stabilize (receiver) countries would similarly mirror either the identity of the sending country (foreign forces sent in to stabilize) or that of the receiver country (local forces engaged by a foreign country to help stabilize).

1.1. Literature

To date, a small body of literature has been emerging in the field of experimental work on delegated decision-making. Following an experiment by Fershtman and Gneezy (2001) on delegation in ultimatum games, a number of studies have investigated the behavioral effects of opportunities to delegate decisions to an agent. A common finding is that, by dividing up a decision, delegation dilutes other-regarding concerns. Hamman, Loewenstein, and Weber (2010) experimentally analyzed allocation decisions taken by agents who compete to be (re)hired by a principal: Principals/owners of capital felt justified in hiring selfish representatives, while these representatives felt justified in acting selfishly in order to conform to their principal's expectations. Fischer, Goerg, and Hamann (2015) replicated and extended the study and found that even in a framed decision context with charity as a recipient and explicit priming to consider the recipient, delegation suppressed monetary transfers. Studying the underlying mechanisms, Bartling and Fischbacher (2011) demonstrated that delegation shifts blame

away from the party that profits from reduced regard to others. Both Oexl and Grossman (2013) and Hill (2015) replicated this effect even in situations where the chosen representative had no power to affect the outcome. Erat (2013) further showed that subjects strategically use this effect when considering lying to others: The more socially unacceptable the lie was, the more likely subjects were to delegate its pronouncement. Similarly, Drugov et al. (2014) demonstrated intermediation to increase corruption.

Trust and reciprocity play important roles in the functioning of economic institutions (Blair & Stout, 2001; Mitchell, 1993) and also for societal cohesion in general (Ashraf, Bohnet, & Piankov, 2006). The idea of bias in favor of one's own group has long been established in literature (Tajfel, 1970; Tajfel & Turner, 1986). According to this "in-group bias," members of a shared group treat each other differently compared to those not part of this in-group. In-group bias has been widely noted experimentally, when groups are defined based on characteristics such as gender (e.g., Bielby, 2000; Rudman & Goodwin, 2004), place of origin/nationality (e.g., Chmura, Goerg, & Weiss, 2016; Glaeser, Laibson, Scheinkman, & Soutter, 2000), and race/ethnic group (e.g., Bertrand & Mullainathan, 2004; Fershtman & Gneezy, 2001). Many of these studies consider in-group bias specifically in the context of bilateral trust or gift-exchange games (Glaeser et al., 2000; Chmura et al., 2016; Fershtman & Gneezy, 2001; see also Johnson & Mislin, 2011 for a meta-analysis of trust games), noting either higher levels of trusting behavior and/or higher levels of reciprocity between players of the same group.

In the South African context, Burns (2006) finds in-group bias among white, but not black, South African subjects in bilateral trust games. Haile, Sadrieh, and Verbon (2008) find an interaction effect of income and racial group on trust among South African respondents, while Zerfu, Zikhali, and Kabenga (2009) note that trust in Africa is influenced by the extent to which people's identities are based on "ethnicity" (identifying primarily with race/tribe rather than country).

Since in- and out-group definitions become unclear when games include three players, we use "own-group" and "other-group" in the remainder of the chapter for clarity.

In many situations where decision-making is delegated (including the examples noted earlier), trust is as important as it is in bilateral interactions. Recently, Maximiano, Sloof, and Sonnemans (2013) and Kvaløy and Luzuriaga (2014) have investigated the impact of delegated decision-making on trust and reciprocity. Kvaløy and Luzuriaga used a trust game set-up; while Maximiano et al. used a gift-exchange game to study worker effort (reciprocity) under different management scenarios. In both cases, the "Client"/"Shareholder" was the owner of the capital; the "Sender"/"Manager" chose a portion of the Client/Shareholder's money to send/pay as a wage to the "Receiver"/"Worker," who decided on the amount of money or effort to return to the Client/Shareholder. The Sender/Manager was paid a fixed fee. In Kvaløy and Luzuriaga, this was done to avoid monetary incentives to care about the Client's returns; while Maximiano et al. included this treatment to study the distinction between the standard gift-exchange game (where the decision-maker (Manager) is the

beneficiary of effort) versus the case where the Shareholder is the sole residual claimant. Their experiment also included a treatment in which the Manager shares in profits. In both cases, the Client/Shareholder took no decisions. Although group identity is not the focus of Kvaløy and Luzuriaga's paper, they do note differences in reciprocity behavior particularly among women when comparing the three-person trust game to the standard bilateral game.[1] Neither paper considers the interaction of group identity (in-group bias) with the delegation decision.

1.2. Research Question

Since existing research has pointed to significant impacts of group identity on bilateral trust, one would expect that own-group bias should also impact trust in scenarios of delegated decision-making. This is the focus of our study: while existing research on delegated decision-making has largely compared the delegation scenario to a no-delegation baseline, we look at the delegation decision itself, and specifically the choice of sender/representative when group identity of the sender might play a role in his/her appointment. If own-group bias does play a role in sender appointment decisions, it is not obvious how this bias would play out: If someone facing a counterpart from the other group is tasked with delegating decisions to either an own- or other-group sender, they might choose an own-group sender due to their own-group bias, or they might choose a sender from the other group to exploit the counterpart's perceived own-group bias strategically. Which of those effects prevails can only be determined empirically. As our research question, we therefore ask how group identity would impact delegation decisions.

To investigate this and find a suitable environment for an empirical test, we turn to South Africa: It has a history of racial separation, formalized through Apartheid ("separate development") laws (1948−1991), which separated education, access to amenities and land ownership for white and non-white South Africans. Many years after the end of Apartheid, racial identity continues to have cultural meaning in South Africa (Seekings, 2008). We explore how this group dynamic affects people's behavior, given their potentially profitable opportunities to exploit each other's biases strategically.

In addition to allowing us better to understand societal issues of trust and ethnic group relations in South Africa, a study in this country offers a natural sample allowing us to examine the role of own-group bias in delegated decision-making interactions involving trust more generally. We therefore conduct an experiment which modifies the original trust game on bilateral interactions[2] to include delegated decision-making and an opportunity for discrimination based on group identity.

We find that the appointment of senders/representatives by owners of capital can be predicted based on differences in expected profit. This is noted for white capital owners: When faced with an other-group receiver, white owners are directionally more likely to select an other-group sender than an own-group sender, anticipating higher profits from this choice. One possible explanation for this profit expectation is that by highlighting the similarity between the sender's

group and that of the receiver, owners attempt to draw their counterpart's focus away from the disparity between their own and their counterpart's group. This use of delegation to shift the focus from the owner's identity in order to maximize profits complements the existing research on delegation. Recall that delegation has been shown to shift people's focus away from their economic counterparty towards his/her representative/agent. We also find, however, that the owners' expectations do not correspond to the actual behavior of receivers. So while rational in terms of owner expectations, their strategy does not ultimately obtain higher profits.

2. EXPERIMENT DESIGN

2.1. Modified Trust Game

In keeping with the previous research cited above, we use the trust game (Berg et al., 1995) as the basis for our experiment. This two-person strategic interaction offers a way to elicit trust preferences effectively. In the one-shot trust game, one player (the "sender") is given a monetary endowment by the experimenter and has to choose which part of this endowment to transfer to a peer player (the "receiver"). The receiver obtains the amount transferred plus twice this amount from the experimenter. The receiver can then choose to return any part of the amount received, and the interaction ends. Where both players are self-interested and rational, the equilibrium outcome of the trust game follows from backward induction: Since the receiver has no incentive to return any amount to the sender, the sender should anticipate this and refrain from transferring anything. At the game's unique subgame-perfect equilibrium, the sender thus earns the initial endowment and the receiver earns nothing.

Any amount that the sender transfers to the receiver can therefore be seen as an indication of trust, while the amount returned by the receiver is an indication of trustworthiness or reciprocity. In general, evidence is found of both trust and reciprocity: A meta-analysis of 162 replications of trust games across 35 countries (Johnson & Mislin, 2011) showed that 50% of the endowment is transferred on average, with 37% of the received amount being returned.[3]

Subjects in our experiment played both this standard trust game and a modified three-player version with the addition of a delegation stage. Our three-player version has three stages: First, the capital owner selects a sender to make decisions about the amount of money to transfer, either from her own-group or from the other group. Second, the sender decides on the amount (out of the owner's endowment of 10) to send to the receiver on behalf of the owner (the owner and receiver have no opportunity to communicate). The amount sent to the receiver is trebled, as in the standard trust game. Finally, the receiver decides on an amount to return to the capital owner out of the trebled amount that he receives through the sender. Please see the Appendix for screenshots showing the decisions made in each stage of the game.

2.2. Treatment Factors: Identity, Knowledge

To compare with existing research, we started all sessions with a standard trust game.

We then introduced the delegation component introduced in "Modified Trust Game," where capital owners were no longer allowed to transfer to receivers directly. Instead, they had to select a "sender" who would transfer money on their behalf, but without any communication.

Both the standard game and the delegation game included the treatment factors set out below:

IDENTITY: Owners (senders in the standard trust game) were paired with receivers of either their own or a different race group (black or white).[4] In the delegation game, owners (who always saw this group information) were given the choice between one black and one white sender and could thus condition on the group parity between themselves and/or between the sender and the receiver. Whether chosen or not, senders received a flat payment equal to the initial endowment, in order to avoid confounding effects of inequality aversion and sender sympathy towards own-group senders.[5]

KNOWLEDGE: We varied whether the receiver was able to see the group of the owner (sender in the standard game). This (in)ability was public knowledge, so the owner could condition his decisions not only on the group of the receiver, but also on whether the receiver was informed or blind as to which type of owner he/she was matched with. Note that the blind condition refers only to the receiver: The sender knew the owner group in this condition.

This gave us four different combinations, detailed in Table 1. All treatment rounds were answered without any feedback about each treatment's results until the end of the experiment, to avoid learning effects which would have jeopardized the mutual independence of observations. Since we were initially worried about the total length of the experiment being too long for continued engagement by participants, we did not include all treatments in the first two sessions. Upon seeing that more tasks could be included within a reasonable time span (~1.5−2 hours), we included all treatments for subsequent sessions in order to increase our participant numbers in each treatment for analysis.[6] The number of participants for each of the treatments is given in Table 1.

We next discuss our experiment protocol in detail.

2.3. Experiment Protocol

Experimental economics is a relatively new field in South Africa, and has not been widely used in the Pretoria/Johannesburg region, where one of the authors is based. This offered a unique opportunity to study a fresh group of people who had never participated in experimental studies before. Yet, it also presented the challenge of not being able to rely on ready-to-use laboratory environments or

Table 1. Treatments and Number of Participants by Session[7].

Session Number	1	2	3	4	5	6	Total Participants
Number of participants	20	20	18	24	26	18	126
Treatments included for group:							
IDENTITY:							
Owner = Receiver	No	No	Yes	Yes	Yes	Yes	86
Owner ≠ Receiver	Yes	Yes	Yes	Yes	Yes	Yes	126
KNOWLEDGE							
Receiver blind to Owner Group	No	Yes	Yes	Yes	Yes	Yes	106
Receiver informed of Owner Group	Yes	No	Yes	Yes	Yes	Yes	106
Treatment combinations							
Owner = Receiver (blind receiver)	No	No	Yes	Yes	Yes	Yes	86
Owner = Receiver (informed receiver)	No	No	Yes	Yes	Yes	Yes	86
Owner ≠ Receiver (blind receiver)	No	Yes	Yes	Yes	Yes	Yes	106
Owner ≠ Receiver (informed receiver)	Yes	No	Yes	Yes	Yes	Yes	106

established recruitment procedures. Both of these were thus developed for the purposes of this study. Research Ethics Clearance was obtained from the Faculty of Economic and Management Sciences at the University of Pretoria, following its established protocols.

Undergraduate economics students were invited using the University of Pretoria online student communication system, "ClickUP," to register for a "decision-making experiment." Interested students followed a link from the "ClickUP" page to register their interest in participating in decision-making experiments. As part of the registration process, ethnic group information was gathered as one of a number of demographic variables. This gave us a subject pool of 412 interested people. Of those who listed their race as either black or white, 58% were black and 42% white. The ethnic group information enabled us to invite balanced groups of 50% black and 50% white respondents to participate in each session, which faithfully approximates the composition of UP's student body (Ouma, 2014, p. 3). We ran six sessions in total. We invited 20 black and 20 white students to each session. We started sessions when we had an equal number of students of each race in the session. Since students in South Africa are frequently not very punctual, we could accept late arrivers until we had an equal ratio, and then start the session. In this way, latecomers who would have changed the race balance of the sessions could be politely turned away on the grounds that the session had already started. Where possible, we started with multiples of six (see Appendix for discussion of the matching process). This, together with some students' not arriving for assigned sessions, resulted in sessions of between 18 and 26 participants, with 126 participants in total. Since there were well over 2000 undergraduate economics students at the time, and since students within the pool of 412 interested in experiments were randomly

selected for specific sessions, we consider it unlikely that many of the students in any individual session knew each other well. Further, by ensuring that the partner(s) in any iteration of the trust game could not be identified with the information given, even if a friend/classmate was in the same session, participants would have no reason to believe they were playing with that person at any point in the experiment.

Upon arrival at the laboratory (set up on campus at the natural sciences graduate computer lab), participants were seated randomly in front of prepared computer terminals. Each terminal was connected via Internet to a web server running a program written in PHP 5.3 (Lerdorf, 1995), exchanging HTML output and web form input with all terminals.

The experiment consisted of four parts. Part 1 was the standard trust game; Part 2 was a number assessment task not used in this study;[8] Part 3 was the modified trust game with delegation; and Part 4 included demographic questions.

Prior to each part of the experiment, participants received printed instructions which were read aloud by a South African instructor.[9] Any screen interaction required of participants was explained and demonstrated live through a projector presentation, giving participants an opportunity to ask any questions afterwards.

Before starting the first stage of the experiment, participants were required to complete a test run, in which they were paired with a randomized computer algorithm to play a one-shot standard trust game. It was made clear to participants that the purpose of this exercise was simply to ensure that they understood how the game would work, and that the payments of the "other player" were random ones generated by the computer. This was reiterated both on screen (where the following text was shown: "Remember that this is a test run. The other person will be simulated by a random computer response, nobody will be paid.") and in the instructions (included in the Appendix for reference). Participants played one unpaid round of the trust game, first in the role of the sender and then in the role of the receiver. In this test round, participants received immediate feedback and had to confirm they had understood the results shown to them. After every participant confirmed and no questions were left unanswered, the on-screen interaction between participants began.

Participants were told that after the test round they would be interacting with other players in the room (although they would only know limited demographic information about each of the players they were paired with). Participants were incentivized to make choices according to their true preferences: all decisions in the game were paid in real money at the end of the experiment. We used a lab currency: lab dollar amounts were divided by two to convert to a payment in South African Rands at the end of the experiment. Expectations of other participants' decisions were elicited, but were not incentivized.

In Sessions 3–6, subjects played eight rounds of the experiment in total. These included four rounds in Part 1 (the standard trust game) and four rounds in Part 3 (the trust game with delegation). In each part, the four rounds consisted of two rounds where the receiver was blind to the owner's group (one of these had owner and receiver from the same group, while the other had owner

and receiver from different groups); and two rounds where the receiver was informed of the owner's group. Within each part, the sequence of these four rounds was randomized. Sessions 1 and 2 included only asymmetric sender-receiver matchings, with Session 1 containing the Informed, and Session 2 the Blind treatment.

In order to maximize the data we collected from each participant, we used the strategy method of experimental interaction (Selten, 1967). In a comparison of strategy method with direct response elicitation, Brandts and Charness (2011) note that they never see a treatment effect with the strategy method that does not exist in the direct response method. Where differences exist, treatment effects appear, if anything, to be understated when the strategy method is used. We therefore suspect that, if anything, our findings will be understated by our choice of method.

With the strategy method, each round in Part 3 included three sets of questions (sender questions, receiver questions, and owner questions), such that players answered each set of questions (and therefore played each available role) once per round. In each role, participants had to state their strategies and expectations for all relevant decisions. These choices are shown in detail in screenshots included in the Appendix and are also discussed below.

In the role of owner, participants had to select a sender, and also had to report the amount they expected each of the two possible senders to transfer to the receiver, as well as the amount they expected the receiver to return in response to each stated transfer. In the role of sender, participants had to state the amount they would transfer in each of two owner−receiver pairings they could possibly be hired into, as well as the amount they expected the receiver to return in response to each transfer. In the role of receiver, participants had to state their return for each of two senders (or owner−sender pairs in the informed scenario) they could possibly confront. In order for participants to play an actual game in which their decisions would be paid at the end, while also avoiding feedback on other participants' decisions during the game, receivers were asked to select a return amount for each of three possible transfers from each sender. Two of these transfer amounts were selected randomly by the program, while the third was the actual amount transferred. The randomization was designed so that receivers always responded to one low transfer (0, 1, or 2), one "average" transfer (3, 4, 5, 6, or 7) and one high transfer (8, 9, or 10).

We do not anticipate that either the change to the session structure after the first sessions (also discussed in "Treatment Factors: Identity, Knowledge") or the decision to use the strategy method of elicitation would have biased responses through experimenter demand characteristics differentially across treatments. As discussed in Zizzo (2010), although order effects can never be eliminated, steps can be taken to minimize these. We took two such steps: First, no payment information (or information about other players' decisions) was given until the very end of the experiment, so that no learning could take place with additional rounds of play. Second, we randomized the sequence of treatment rounds in all sessions, including those where more treatments were

included. In this way, if cooperation were to increase with additional rounds, this benefit would be shared across treatments since different players would have had more/fewer priming rounds before facing each treatment condition. If anything, increasing cooperation would lead to smaller differences between treatments, such that we might understate any treatment effects through the inclusion of multiple roles. Another possible mechanism for increased cooperation with the playing of multiple roles is that risk averse players might play more cooperatively to mitigate the risk of being in the role that is most often hurt in trust interactions. By paying players for all decisions made in the game, we remove this incentive for different behavior in the strategy method.

Since the standard trust game was always played before the delegation game, it is possible that some order bias might exist were one directly to compare the findings between the two games. Our goal, however, was not to do this. We included the standard trust game primarily to see how our findings compared to earlier research on trust and race in South Africa (Burns, 2006).

Our approach to providing race information to participants built on work by Bornhorst, Ichino, Kirchkamp, Schlag, and Winter (2010) who conveyed nationality as one piece of information among others that could be assumed (and later shown) to be inconsequential. Similarly, we presented virtual information cards (see screenshots in the Appendix for an illustration) containing racial group and three decoy demographics that were designed to be irrelevant:[10] (1) whether a person's age was below or above 30 years; (2) whether a person's parents had an age difference of more or less than three years; (3) whether a person's parents, brothers, and sisters numbered even or odd. While (1) was conceivably relevant, but invariant in our subject pool, (2) and (3) were supposed to have some air of relevance while not influencing decisions.[11] Previous experiments used racial group-specific surnames (van der Merwe & Burns, 2008), but this would have been prohibitively difficult in the campus setting where anonymity had to be guaranteed despite conveying identity information.

After the experiment, participants answered demographic questions and were then presented with the outcomes of all previous decisions and with their calculated final pay-off: All tasks were paid based on actual decisions made by the participants and the players with whom they were paired at different points in the experiment. After checking and confirming this feedback, participants were sent out of the lab one at a time to receive their payment in cash.

All sessions lasted one to two hours with average payments of 111.57 ZAR – an equivalent of about $9 USD, or 1.5 paid hours for an assistant lecturer (part-time student assistant) in the Economics department at the University of Pretoria, around the time of the experiment.

In keeping with the ethics requirements of the University of Pretoria, subjects were assured that all their responses would be anonymous, participation was voluntary and terminable at any time, and no names or personal identifiers were recorded with the data.

3. HYPOTHESES

We broadly consider two questions which guide the development of our hypotheses: (1) How does behavior (in the role of owner, sender, and receiver) change depending on the group(s) of other players? (2) How does the presence or absence of information about players' group impact behavior? Our specific hypotheses consider these questions in the three stages of the delegation game following the pairing of owner and receiver: owner's decision to appoint a sender, sender's transfer decision, and receiver's return decision. Since participants play all three roles (owner, sender, and receiver) in the course of the game, we assume that they are aware of the incentives of those in other roles. We therefore start with the receiver decision and assume that, in the role of sender, players would consider the expected decision of the receiver; while in the role of owner, players would consider the expected decision of the sender, which in turn would be based on the expected decision of the receiver.

Similarly, when the receiver is blind to the group identity of the owner, we assume that the receiver would consider the expected sender appointment decision of the owner in deciding on a return amount. The pairing of own or other-group owners and receivers was expected to impact decisions when there is a positive probability of own-group bias.

3.1. Return Decision by Receiver

We take the case where the receiver knows the owner group, and the owner group is the same as that of the receiver, as our benchmark case. Here, we expect that receivers would transfer "average" returns (Johnson and Mislin's (2011) meta-analysis of trust games finds an average of ~37% returned) to a same-group owner. Since the receiver has no way of punishing or rewarding the sender (given the fixed payment of the sender), *ceteris paribus* the group of the sender should not impact the return decision of the receiver. Where the receiver is aware of the owner group, we expect returns to be conditioned on the owner's group rather than the sender's group.

When the receiver knows that the owner group is different from her own one, a biased receiver from the owner's own group would return an "average" amount, as above, while a biased receiver from the owner's other-group might return nothing (or at least, would return an amount strictly less than the "average" amount). Since there is no repeat play of the game and hence no strategic benefit to sharing for the receiver, a biased receiver has no incentive to share his earnings with an other-group owner. With only a weak assumption of a nonzero probability of own-group bias among receivers, we expect that the return amount to a receiver-group owner would be strictly greater than the return amount to a nonreceiver-group owner in this scenario.

We next consider the "blind" scenario, in which the owner still chooses the sender by group, but the receiver does not know the group of the owner. The receiver might interpret the owner's sender selection decision as a signal of owner group. Here the receiver would make assumptions about the owner group

based on the believed probability of owners selecting the same – or other – group senders in the scenario where the owner is aware that the receiver does not know the owner's group (this is set out in detail in "Experiment Protocol"). The selection of a receiver-group sender is expected to be dominant for the owner in this "blind" scenario: Where the receiver and owner are of the same group, the owner is expected to choose an own-group (=receiver group) sender; and where the owner and receiver groups differ, the owner is expected to choose a receiver-group sender.

Deviations from receiver-group sender might arise when an other-group owner believes that a sender of the owner's own group (instead of the receiver's own group) might better represent his interests (e.g. by transferring less of his capital to a potentially biased receiver). The receiver might then interpret the nonreceiver-group sender as a signal of a biased other-group owner. As such, even an unbiased receiver might be tempted to reduce the amount he would otherwise return.

Since no clear signal of owner group can be obtained from a receiver-group sender choice, we assume that the receiver returns an "average" amount in this situation. Where a nonreceiver-group sender is encountered in the "blind" scenario, we anticipate a strictly lower receiver return.

Receiver return predictions based on the outline above are summarized in Table 2.

We therefore posit the following hypotheses:

H1: Amounts returned when $S = R$ *(informed)* $\approx$ amounts returned when $S \neq R$ *(informed)*

H2: Amounts returned when $S = R$ *(blind)* $>$ amounts returned when $S \neq R$ *(blind)*

H3: Amounts returned when $R = O$ *(informed)* $>$ amounts returned when $R \neq O$ *(informed)*

Table 2. Expected Receiver Decisions.

Variation	Explanation	Hypothesized Return Transfer
$S = R$ (Informed)	Sender same group as receiver	Influenced by owner group: no impact of sender group
$S \neq R$ (Informed)	Sender different group to receiver	Influenced by owner group: no impact of sender group
$R = O$ (Informed)	Receiver same group as owner	Average
$R \neq O$ (Informed)	Receiver different group to owner	Below average
$S = R$ (Blind)	Sender same group as receiver	Average
$S \neq R$ (Blind)	Sender different group to receiver	Below average

3.2. Transfer Decision by Sender

Recall that the sender is paid a fixed fee for his/her decision-making. Further, the nonappointed sender receives the same fixed fee. The sender therefore does not learn whether or not he was selected during the trust game. Since the sender receives no benefit from being selected, she should have no incentive to repay the owner for her selection. We therefore assume that a neutral sender would favor a fair (~50/50) division of profits. We consider two possible motives for the sender's choice:

(1) Preference for efficiency: any decision by the sender to transfer money increases the total resources to be shared between owner and receiver (since any money transferred is trebled before reaching the receiver).
(2) Equity concerns: the sender's decision together with the receiver's return decision will determine the proportion of the total available funds accruing to owner and receiver. The sender must therefore take the (assumed) return decision of the receiver into account when deciding on a transfer amount.

An unbiased sender would be expected to target both efficiency and equity by transferring something in the region of the 50% average seen in trust games (Johnson & Mislin, 2011).[12] We summarize the eight scenarios for this transfer (four in which the receiver knows the group of the owner ("informed"), and four in which the receiver is "blind" to owner group) in Table 3. Recall that the sender is always aware of the group of both owner and receiver.

We would expect to see higher transfers by senders to sender-group receivers, although the difference would be slight when receiver and owner are of the same group. Aligning with this predicted outcome, the literature on group identity predicts that each member of a group prefers to trust other members of this group over members of other groups (own-group favoritism).

Considering the possibility for a biased sender to discriminate against both other players by limiting the transfer amount (limiting efficiency, as noted earlier), we would anticipate that lower amounts would be transferred when the sender group is different from the group of both the receiver and the owner.

We therefore propose:

H4: Amounts transferred when $S = R >$ amounts transferred when $S \neq R$, particularly when $O \neq R$.

3.3. Sender Appointment by Owner

The owner faces four scenarios in which she has to choose between an own-group and other-group sender. As before, we assume that the owner is aware of the likely transfer decision of the sender, and of the likely return decision of the receiver (as discussed in "Return Decision by Receiver" and "Transfer Decision by Sender").

The predicted outcomes are summarized in Table 4.

When the owner group is known to the receiver, we expect the owner to show some favoritism towards owner-group senders (based on possible lower transfers from other-group senders) when facing an own-group receiver; and we

Table 3. Sender Transfers.

Variation	Explanation	Hypothesized Transfer	Reason
$O = R = S$ (Informed)	Receiver same group as owner; same-group sender	Average	No discrimination possible
$O = R = S$ (Blind)	Receiver same group as owner; same-group sender	Average	No discrimination: sender knows that receiver can't predict owner group with a same-group sender
$O = R \neq S$ (Informed)	Receiver same group as owner; other-group sender	(Slightly) below average	Biased sender could limit efficiency: lower transfer would discriminate against both other players
$O = R \neq S$ (Blind)	Receiver same group as owner; other-group sender	(Slightly) below average	Biased sender might limit efficiency. If receiver is expected to misinterpret sender group as signal of biased out-group owner, nonbiased sender might limit transfer because of expected low return
$O = S \neq R$ (Informed)	Receiver different group to owner; owner-group sender	Below average	Biased sender would transfer below average (possibly 0). Nonbiased sender would transfer average
$O = S \neq R$ (Blind)	Receiver different group to owner; owner-group sender	Below average	Biased sender would transfer less than average (possibly 0). If receiver is expected to misinterpret sender group as signal of biased out-group owner, nonbiased sender might limit transfer because of expected low return
$R = S \neq O$ (Informed)	Receiver different group to owner; receiver-group sender	Above average	Biased sender would transfer above average (possibly 100%). Nonbiased sender would transfer average
$R = S \neq O$ (Blind)	Receiver different group to owner; receiver-group sender	Above average	Biased sender would transfer above average. Nonbiased sender would transfer average

expect strict preference for an owner-group sender in facing an other-group receiver. When the owner group is not known to the receiver, we expect an owner to prefer a sender from the same group as the receiver (and himself) if owner and receiver share a group: A nonreceiver-group sender risks sending a false signal of a biased other-group owner to the receiver. Similarly, we anticipate a preference for receiver-group senders (now nonowner group) where the owner and receiver group are different. The most notable difference is seen when we compare the cases of informed and blind receiver where receiver and owner group differs: here, we expect strict preference for a nonreceiver-group sender when facing an informed receiver, but strict preference for a receiver-group sender when facing a blind receiver. That is, we expect the owner to use the receiver's blindness to her strategic advantage.

Table 4. Likelihood of Selecting an Owner-group Sender.

Variation	Explanation	Hypothesized Preference	Reason
$R = O$ (Informed)	Receiver same group as owner	Some preference for owner (=receiver) group sender	Risk of lower transfer from biased other-group sender; preference for own-group sender (in-group bias) also possible
$R \neq O$ (Informed)	Receiver different group to owner	Strict preference for owner (=nonreceiver) group sender	Biased receiver-group sender might transfer all to biased other-group receiver who might return low amount. Biased owner-group sender might send low amount and biased receiver might return low amount, but the balance of the endowment would be higher due to the low transfer by the biased sender
$R = O$ (Blind)	Receiver same group as owner	Strict preference for owner (=receiver) group sender	Average transfer and return expected with receiver-group sender. Nonreceiver sender could decrease return from receiver if this is misinterpreted as signaling a biased other-group owner; Biased other-group sender might select low transfer to discriminate against both receiver and owner
R ≠ O (Blind)	Receiver different group to owner	Strict preference for other (=receiver) group sender	Although biased owner-group sender might transfer less (keeping more for owner), biased receiver would return little and even nonbiased receiver might punish perceived bias of owner (recall from "Return Decision by Receiver" that a nonreceiver-group sender could be interpreted as a signal of a biased owner); biased receiver-group sender might transfer more, but receiver should offer average return to a receiver-group sender where owner group is unknown, resulting in higher total returns to owner

We hypothesize:
H5: Probability of selecting $S = O$ when $R \neq O$ *(informed)* > probability of selecting $S = O$ when $R \neq O$ *(blind)*.

4. RESULTS

We conducted six experimental sessions at the University of Pretoria (Hatfield campus) between March 19, 2015 and April 23, 2015.

4.1. Demographics

Of the 126 subjects participating in the experimental sessions, 67 were female (53 %), and 65 were black (52 %). Participants ranged in age from 18 to 26 years, with an average of 19.3 years. Excepting two, all subjects were in the first or second year of their studies. Demographics by race group are reported in Table A1 in the Appendix.

4.2. *Descriptives: Trust and Race in South Africa*

To contextualize our study within previous research, we first compare findings from our standard trust game with those from existing studies conducted in South Africa. Recall that, in earlier research using the standard trust game (Burns, 2006), whites showed insider bias[13] (transferring more to white partners than to nonwhite partners). Similar insider bias was shown by white respondents in a dictator game in South Africa (van der Merwe & Burns, 2008). To compare this with our data qualitatively, we first inspect interactions between the sender and receiver race group in the different constellations of our standard trust game. Note that since we used the strategy method of elicitation, in which actual transfers and returns were not revealed until all rounds had been concluded, we align actual transfers with associated expected returns; and actual returns with associated expected transfers.

Each of the eight pairs of bars in Fig. 1 compares behavior across the race group of the respective receiver in the right panel, and owner/sender in the left panel. Fig. 1 reveals own-group bias among white, but not black senders: They transfer significantly higher amounts (7% on average, $p = 0.09$) to white than to black receivers and expect to receive significantly higher returns (16% on average, $p = 0.02$) from a white than from a black receiver.[14]

Burns' (2006) research with high-school students found general discrimination against black respondents (even black senders transferred more to white than to black partners). Neither later research on dictator games (van der Merwe & Burns, 2008) nor our own data (Wilcoxon signed-rank test, $z = -0.73$, $p = 0.47$, $n = 126$) confirm this pattern of generally lower transfers to black receivers. Burns attributed the discrimination against black students to mistaken stereotypes – black receivers were expected to return a lower proportion of the received amount on average – but our data do not confirm that black receivers are expected to return significantly less than white receivers (Wilcoxon signed-rank test, $z = -0.92$, $p = 0.36$, $n = 126$).

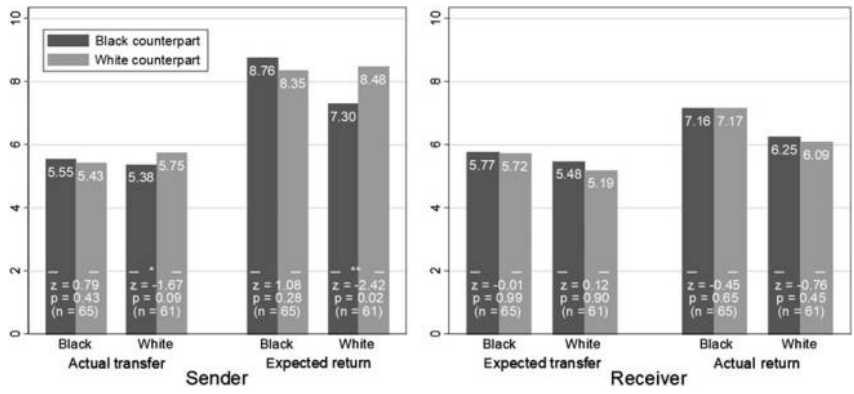

Fig. 1. Decisions by Group Identity in the Standard Trust Game. *Note*: Reported statistics are for Wilcoxon signed-rank tests.

These differences may be explained either by the time that passed between the studies (apartheid is one more decade away) or by the different study contexts: While Burns (2006) used high-school students, where multiple schools were needed to achieve sufficient racial heterogeneity, both van der Merwe & Burns (2008) and our own study were conducted with university student participants of considerable racial heterogeneity. As noted in the experiment protocol, our decision to use the strategy method could have diluted results, such that results that appear only borderline significant might have shown up as clearly significant had we used a direct approach.

4.3. Amounts Returned by Receivers

Some systematic differences exist between owner expectations and actual behavior. We therefore consider for each set of hypotheses, first, whether owner expectations align with the expectations model proposed in "Hypotheses"; and second, whether actual behavior aligns with our predictions. Where owner expectations differ from our predictions in "Hypotheses", we can then see whether this is because owners correctly anticipated behavioral deviations from our model; or whether owners' expectations are incorrect.

Our first three hypotheses proposed that when sender and receiver shared a group, we would see no impact of sender group on amounts returned when owner group is known (*H1*); but that amounts returned would be higher to a receiver-group sender when owner group is not known (*H2*). We further proposed in our *H3* that when owner group is known, amounts returned to a receiver-group owner would be higher than amounts returned to a nonreceiver-group owner.

Fig. 2 shows that owner expectations do not align well with our predictions. In particular, contrary to *H1*, owners expected receivers to condition on the parity/disparity between sender and receiver group, particularly when the receiver knew the owner's group. In terms of *H2*, although directionally owners also expected receivers who are blind to the owner's group to condition on the sender group, this effect is not significant. For *H3*, there was no expectation of significantly

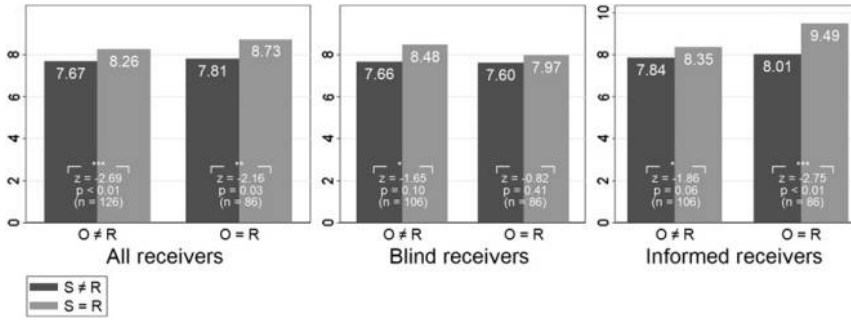

Fig. 2. Owner Expectations of Returns by Receiver: Impact of Group (Dis)parity.
Note: Reported statistics are for Wilcoxon signed-rank tests.

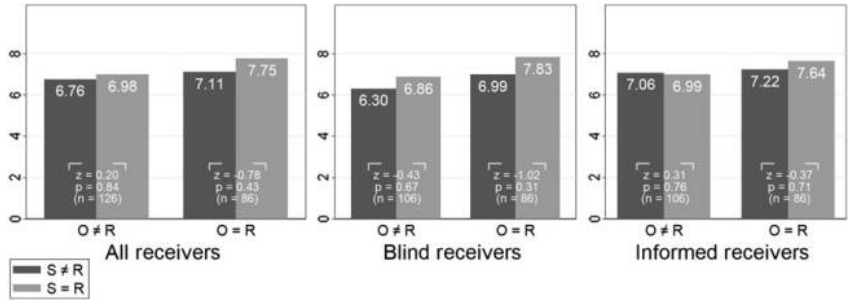

Fig. 3. Receiver Behavior: Impact of Group (Dis)parity on Return Amount. *Note*: Reported statistics are for Wilcoxon signed-rank tests.

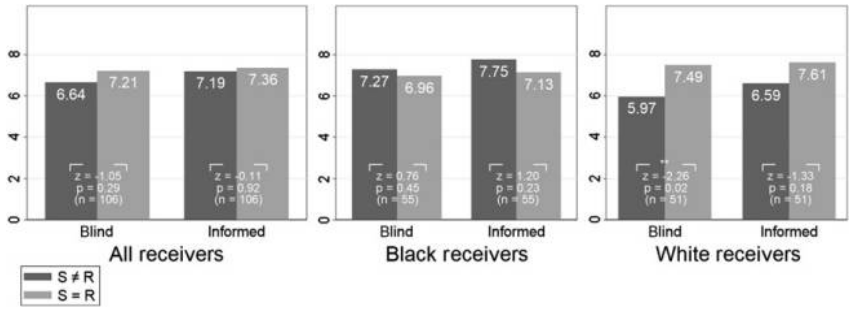

Fig. 4. Receiver Behavior: Amounts Returned from Blind and Informed Receivers by Group (Dis)parity. *Note*: Reported statistics are for Wilcoxon signed-rank tests.

higher amounts returned to receiver-group owners when the receiver and sender group was the same (Wilcoxon rank-sum test: $z = -1.27$, $p = 0.20$, $n = 106$).[15]

We next turn to actual behavior of receivers.

Contrary to owner expectations, actual behavior, shown in Fig. 3, aligned with our *H1* for informed receivers: No significant differences in amounts returned were seen between different sender groups, with either other-group or own-group owner[1617]. However, *H3* is not confirmed: Although amounts returned by informed receivers (receivers who are aware of the owner's group) were slightly higher when the owner was of the same group as the receiver, this difference is not significant (Wilcoxon rank-sum test: $z = -0.84$, $p = 0.40$, $n = 106$).[18]

Testing *H2* against actual receiver behavior, Fig. 4 shows that, with blind receivers, amounts returned were slightly higher on average with receiver-group senders versus other-group senders. As an exploratory analysis, we consider black and white receivers separately. Here, we see that this difference was driven entirely by white receivers, for whom amounts returned were significantly higher with receiver-group senders. Recall that a nonreceiver-group sender was

expected to indicate a biased nonreceiver-group owner, suggesting that white receivers indeed discriminate against suspected (biased) other-group owners.[19]

To summarize, our first three hypotheses met with mixed results:

> *H1:* Amounts returned when $S = R$ *(informed)* $\approx$ amounts returned when $S \neq R$ *(informed)*.

Supported (we cannot reject the hypothesis of no significant different)

> *H2:* Amounts returned when $S = R$ *(blind)* > amounts returned when $S \neq R$ *(blind)*.

Supported only for whites (null hypothesis of no significant difference can only be rejected for whites)

> *H3:* Amounts returned when $R = O$ *(informed)* > amounts returned when $R \neq O$ *(informed)*.

Not supported (null hypothesis of no significant difference cannot be rejected)

4.4. Sender Transfers

H4 considered sender transfers, and whether these would be influenced by racial (dis)parity between owner and receiver and/or between sender and receiver. We anticipated higher transfers to sender-group receivers than to other-group receivers. This difference was expected to be slight when the owner and receiver group was the same, but significant when they were different. Again we start with the owner expectations, shown in Fig. 5.

As Fig. 5 shows, the owners' expectations aligned well with our *H4*: senders were expected to transfer more to receivers of the same group as the sender than to receivers of a difference group. Again, as exploratory analysis we consider black and white owners separately and note that, for black owners, this was only expected to be the case when owner and receiver group were different.[20]

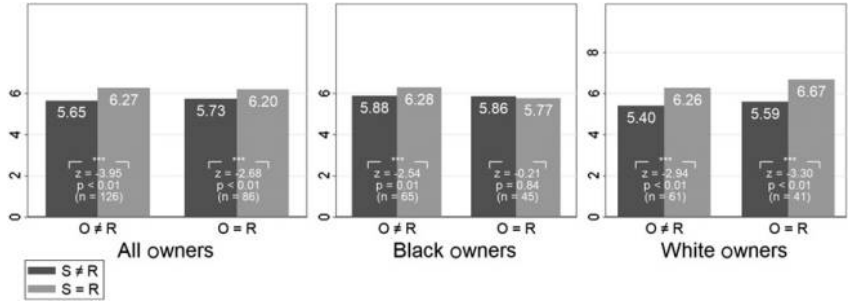

Fig. 5. Owner Expectations: Sender Transfers by Group (Dis)parity. *Note*: Reported statistics are for Wilcoxon signed-rank tests.

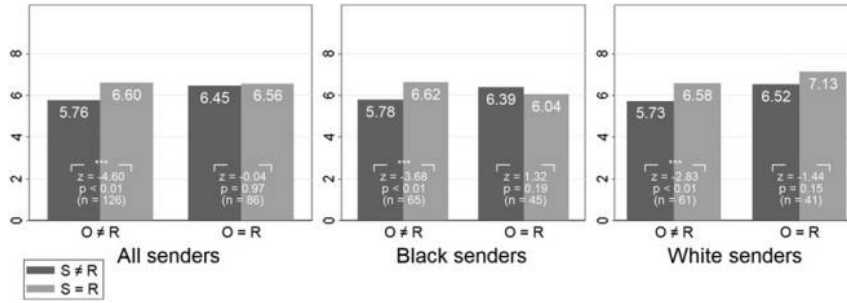

Fig. 6. Sender Behavior: Transfers by Group (Dis)parity. *Note*: Reported statistics are for Wilcoxon signed-rank tests.

We next consider actual sender behavior.

Fig. 6 again includes an exploratory consideration of black and white senders separately. The figure shows that actual transfers differed by (dis)parity between sender and receiver group, but only when owner and receiver group is different: Significant differences were seen between same and other-group senders when owner and receiver group differed; but no significant difference was seen (for either group) when owner and receiver group were the same. This aligns well with the predictions of our *H4*: senders discriminate against other-group receivers more when the owner is from the sender's own group. Interestingly, when all three player's groups align (i.e., when owner group = sender group = receiver group), we note higher transfers when all players are white than when all players are black.[21]

> *H4:* Amounts transferred when $S = R >$ amounts transferred when $S \neq R$, particularly when $O \neq R$: *Supported* (we cannot reject the null hypothesis)

4.5. *Sender Selection*

Our final hypothesis predicted differences in sender selection based on racial (dis)parity between owner and receiver. Since this is an owner decision, we show only actual sender selection in Fig. 7.

Recall that *H5* predicted that, when the owner and receiver group are different, owners are more likely to select an owner-group sender when the receiver is informed of the owner's group than when the receiver is blind to the owner's group. Specifically, we anticipated that the owner should always appoint an own-group sender in the former condition, but should never appoint an own-group sender in the latter. Yet, the former condition sees the lowest likelihood of own-group sender selection. This is likely related to owner expectations: recall that owners expected higher transfers and returns when sender and receiver group are the same, such that selecting a receiver-group sender would result in higher profits for the owner. Perhaps, the owner hopes to use the other-group sender to draw focus away from the known group discrepancy between herself and the other-group receiver. By presenting themselves as "good-willed" and other-regarding, owners might hope to elicit higher returns.[22]

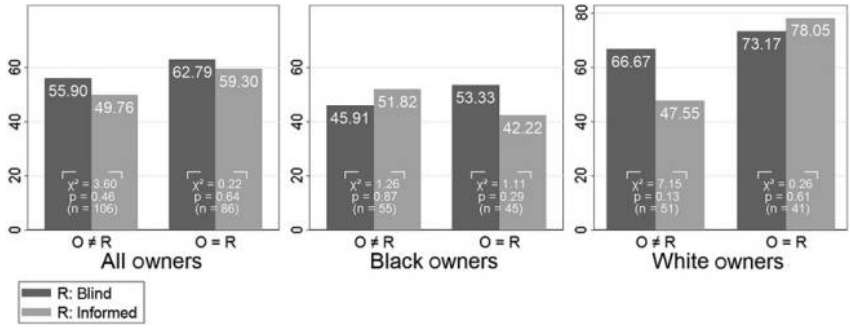

Fig. 7. Sender Delection: Proportion of Owners Selecting an Owner-Group Sender.

H5: Probability of selecting $S = O$ when $R \neq O$ *(informed)* > probability of selecting $S = O$ when $R \neq O$ *(blind)*: *Rejected*

We summarize our main findings over our five hypotheses in Table 5. Although our initial hypotheses did not separate black and white players, we explored differences between these groups in seeking explanations for our overall findings.

The owners' sender choices do not align well with our predictions, particularly for white owners. While our hypothesis was rejected for both black and white owners, we noted a surprising directional preference for receiver-group senders by white owners when the owner group was known to the receiver. Although this finding is not statistically significant, we note again that the strategy method can weaken results. We therefore investigate it as an interesting directional trend, and note that it appears to be related to the owners' (incorrect) expectations of receiver behavior: recall from Fig. 2 that owners expected higher returns to receiver-group senders, even when the owner group was known to the receiver (although this was not reflected in actual returns). White owners also expected higher transfers whenever sender and receiver group are the same (shown in Fig. 5). This expectation of higher transfers and returns with a receiver-group sender offers an explanation for white owners showing some preference for black senders when confronting black receivers.[23]

5. CONCLUDING REMARKS

This chapter investigated the role of group identity on delegated decision-making, and particularly on the selection of a sender to whom the investment decision would be delegated. We conducted a modified trust game experiment where capital ownership was divorced from decision-making, and where the capital owner was asked to select a representative decision-maker (sender) where the potential senders either shared or did not share a group identity with the owner. The group identities (operationalized as race group identities in South Africa, where race remains highly salient) of owner, sender, and receiver were

Table 5. Summary of Results by Hypothesis.

		Hypotheses	Exploratory: Black	Exploratory: White
H1	An informed (knows the owner's group) receiver will not condition return amount on the group of the sender (H_1: expect no significant difference)	*Supported:* we cannot reject the hypothesis of no significant difference ($p = 0.92$)	*Supported* ($p = 0.23$)	*Supported* ($p = 0.18$)
H2	A blind (does not know the owner's group) receiver will return more to an own-group sender than to an other-group sender (H_1: expect significant difference)	*Rejected:* we cannot reject the H_0 of no significant difference ($p = 0.29$)	*Rejected* ($p = 0.45$)	*Supported* (we reject H_0) ($p = 0.02$**)
H3	An informed receiver will return more to an own-group owner than to an other-group owner (H_1: expect significant difference)	*Rejected:* we cannot reject the H_0 of no significant difference ($p = 0.4$)	*Rejected* ($p = 0.91$)	*Rejected* ($p = 0.15$)
H4	A sender will transfer more to a sender-group receiver than to a non-sender-group receiver, particularly where owner and receiver group differs (H_1: expect significant difference)	*Supported:* we reject the H_0 of no significant difference where $O \neq R$ ($p < 0.01$***)	*Supported* where $O \neq$ R ($p < 0.01$***)	*Supported* where $O \neq$ R ($p < 0.01$***)
H5	Where owner and receiver group differs, owners facing informed receivers are more likely than owners facing blind receivers to select an owner-group sender (H_1: expect significant difference)	*Rejected:* we cannot reject the hypothesis of no significant difference ($p = 0.46$)	*Rejected* ($p = 0.87$)	*Rejected* ($p = 0.13$) (Informed receivers are directionally less likely to select an owner-group sender)

Notes: ***$p < 0.01$, **$p < 0.05$.

systematically varied for players in all three roles in order to understand the impact of group identity in a delegated decision-making context.

We predicted the behavior of receivers, senders, and owners where we assumed that all players take into account the impact of possible own-group biased behavior of other players. Sender transfers and receiver returns were directionally in line with most of our predictions, driven largely by the behavior of white respondents: senders transferred more to own-group receivers when the owner was other group; and when receivers were not aware of the owner group, white receivers returned more when facing a sender of their own group.

Sender appointment decisions, however, differed systematically from our predictions. In particular, while we expected that concealing the identity of the owner from the receiver might lead to increased selection of other-group senders when facing other-group receivers (indeed, this was the only time when selecting an other-group sender was expected to be a dominant strategy), we did not see this in our data. Instead, (white) owners were directionally (although not significantly) more likely to select an other-group sender when facing an other-group receiver to whom their own (white) identity was known. We note that owners expect higher profits from appointing receiver-group senders. This ties in with a body of literature suggesting that delegation can be used to shift the focus away from the opposing player to his/her representative: White owners might be hoping to draw attention away from their own (different) race group.

We end by noting some interesting directions for further research on the impact of group identity on trust in delegated decision situations:

Since actual returns do not conform to owner expectations in our experiment, the apparent strategic use of group identity by white owners does not increase profits. This raises the question whether different results might be seen in an environment where learning is permitted. Replications with repeated interactions with feedback could shed light on this question.

Our findings suggest that, where group identity is salient, this might impact delegation decisions. Some of our findings are, however, either borderline significant or only directional. A first step in future research could involve a similar experiment where direct responses are used instead of the strategy method employed in this experiment. Since the strategy method can weaken results by diluting the salience of individual decisions, some of our results with borderline significance might become significant with this method.

The higher prevalence of own-group bias in white versus black respondents suggests that the impact of group identity varies by cultural group. Research in different cultural contexts (including a less homogenous group with varied ages and backgrounds in South Africa) could confirm whether the biases we see in white respondents are common or exceptional.

NOTES

1. Kvaløy and Luzuriaga's paper compares sending and receiving behavior where senders transfer their own money (as in the standard trust game set-up) to behavior where

senders transfer money belonging to a separate player (capital owner). They note that women return less to the capital owner when the transfer amount was decided by another player and more when the capital owner was also the decision-maker for the transfer amount. This finding suggests that reciprocity motivates women's return decisions.

2. See the seminal contribution of Berg, Dickhaut, and McCabe (1995), and, for a meta-analysis of experimental trust game studies, Johnson and Mislin (2011).

3. Variability of transfers and returns is fairly high: The standard deviation for transfers is 12%, while that for returns is 11%.

4. Where the receiver does not know the race of the principal, the balanced (50% white, 50% black) session set-up would make any guess about the principal race random.

5. This flat payment to the sender/intermediary also aligns with the approch used by Kvaløy and Luzuriaga (2014) and Maximiano et al. (2013).

6. Where regressions were run for robustness checks, we included a dummy variable for sessions 1 and 2, since these differed in terms of treatments. This dummy variable was not significant in any of our regressions.

7. Since all players played all three roles in the delegation game (and both roles in the standard trust game), participant numbers are the same for owners, senders, and receivers.

8. This task involved asking respondents to guess transfer and return amounts from players in the standard trust game. Any biasing impact that such a task might have on decisions made in the delegation game (Step 3) should be consistent across treatments within this game, and should therefore not impact our analysis of differences between the different treatments in this game.

9. The same (white) instructor was used across all sessions to avoid possible confounds from instructor race. The instructor was a PhD student not usually based on campus, and so was unknown to students. The instructor's preferences/beliefs are therefore not believed to have led to experimenter demand effects.

10. To confirm that decoys were inconsequential, we later regressed transfer on each of the presented demographics in interaction with the corresponding demographic of the decision-maker, and on a dummy for whether the decision was made by the sender or the owner. This Tobit regression (with standard errors clustered on individuals) revealed significant effects of opponent's race ($p < 0.01$), interaction between decision-maker's and counterpart's race ($p < 0.01$), and sender/owner dummy ($p < 0.02$); no other variable had significant effect ($p > 0.12$). The entire model was significant at $p < 0.02$.

11. Upon seeing indications that they might do that (some subjects reported that " if your parents' age difference is more than three years, I assumed that the person came from a wealthy family no matter what" and misunderstandings such as " I mostly looked at the even or odd parents as that plays a part in the upbringing of a person."), we did not present decoys for the last two sessions, by graying out the decoy fields of the information cards. To ensure this did not alter results, we later regressed transfers on a dummy for whether decoys were present or absent, and an owner/sender dummy. While the latter was significant in a Tobit regression ($p < 0.01$), as well as 5 of the 10 possible transfer levels in a multinomial logistic regression ($p < 0.10$), the former was insignificant throughout ($p = 0.13$ in the Tobit regression, and $0.73 > p > 0.11$ in the multinomial logistic regressions on each possible transfer amount).

12. An average transfer of ~50% being trebled, and an average return of ~37% of the trebled amount, together result in an approximately equal split between principal, receiver, and agent (given the agent's fixed fee of 10).

13. Burns also found that black players transferred more to white players than to fellow black players, a result which is not replicated in our data.

14. As an additional robustness test, we regressed the sender's transfers and expected returns and the receiver's return amounts and expected transfers on the race of the other player and the interaction between the two players' races. Neither of these predictors is statistically significant.

15. As an additional robustness check, we regressed expected returns on the variables portrayed in the graph. Only the informed receiver when $O = R$ and $S = R$ has a significant positive coefficient ($p = 0.02$).

16. When both groups ($O = R$ and $O \neq R$) are combined, a Wilcoxon signed-rank test comparing returns with a receiver-group agent to returns with a nonreceiver group agent indicates no significant difference ($z = 0.11$, $p = 0.92$).

17. As an additional robustness check, we confirm these findings using regressions. None of the predictors shown in the graph are significant in a regression considering receiver return amounts.

18. The directional effect is driven by white receivers: For black receivers, the Wilcoxon rank-sum test gives $z = 0.12$, $p = 0.91$ ($n = 55$), while for white receivers $z = -1.44$, $p = 0.15$ ($n = 51$).

19. As a robustness check, we conduct a regression using the same predictor variables as shown in Fig. 4. This gives similar results: only blind white receivers where $S \neq R$ show significance at $p = 0.10$.

20. As a robustness check, we run a regression with the same predictor variables seen in Fig. 5. We see two significant predictors: black owners where $O \neq R$ and $S = R$ has a positive coefficient ($p < 0.01$); and white owners where $O = R$ and $S = R$ has a positive coefficient ($p = 0.04$). White owners where $O \neq R$ and $S = R$ is only significant at $p = 0.10$.

21. As a robustness check, we run a regression replicating Fig. 6 predictors. We note that sender transfers are higher (positive coefficients) for black senders where $S \neq R$ and $O = R$ ($p = 0.08$) and where $S = R$ and $O \neq R$ ($p < 0.01$); and for white senders where $S = R$ and $O = R$ ($p < 0.01$).

22. Regression analysis confirms that the only time owner-group agents are not significantly more likely to be selected by white owners is where $O \neq R$ and the receiver is informed of the principal race.

23. Where expected transfers and returns of selected versus nonselected agents are examined, we see that selected agents are expected to generate higher profit for the principal than nonselected agents.

24. As a further robustness check, we add a dummy variable, "session12", to test for any impact of slight differences in the set-up of the first two sessions (as detailed in Table 1). This dummy is not significant.

ACKNOWLEDGMENT

We thank computer lab manager Mishwell Mathebula for his repeated technological support in conducting our experiment, and gratefully acknowledge funding by the Max Planck Society. We also thank the editor and four anonymous referees, as well as Amalia Álvarez Benjumea, Brian Cooper, Christoph Engel, Lars Freund, Oliver Kirchkamp, Don Ross and Alex Zimper for helpful feedback.

REFERENCES

Ashraf, N., Bohnet, I., & Piankov, N. (2006). Decomposing trust and trustworthiness. *Experimental Economics*, *9*(3), 193–208.

Bartling, B., & Fischbacher, U. (2011). Shifting the blame: On delegation and responsibility. *Review of Economic Studies*, *79*(1), 67–87.

Berg, J., Dickhaut, J., & McCabe, K. (1995). Trust, reciprocity, and social history. *Games and Economic Behavior*, *10*(1), 122–142.

Bertrand, M., & Mullainathan, S. (2004). Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. *The American Economic Review*, *94*(4), 991–1013.

Bielby, W. T. (2000). Minimizing workplace gender and racial bias. *Contemporary Sociology*, *29*(1), 120–129.

Blair, M. M., & Stout, L. A. (2001). Trust, trustworthiness and the behavioural foundations of corporate law. *University of Pensylvania Law Review*, *149*(6), 1735–1810.

Bornhorst, F., Ichino, A., Kirchkamp, O., Schlag, K. H., & Winter, E. (2010). Similarities and differences when building trust: The role of cultures. *Experimental Economics*, *13*(3), 260–283.

Brandts, J., & Charness, G. (2011). The strategy versus the direct-response method: A first survey of experimental comparisons. *Experimental Economics*, *14*(3), 375–398.

Burns, J. (2006). Racial stereotypes, stigma and trust in post-apartheid South Africa. *Economic Modelling*, *23*(5), 805–821.

Chmura, T., Goerg, S. J., & Weiss, P. (2016). Natural groups and economic characteristics as driving forces of wage discrimination. *European Economic Review*, *90*, 178–200.

Drugov, M., Hamman, J., & Serra, D. (2014). Intermediaries in corruption: An experiment. *Experimental Economics*, *17*(1), 78–99.

Erat, S. (2013). Avoiding lying: The case of delegated deception. *Journal of Economic Behavior & Organization*, *93*, 273–278.

Fershtman, C., & Gneezy, U. (2001). Discrimination in a segmented society: An experimental approach. *The Quarterly Journal of Economics*, *116*(1), 351–377.

Fischer, S., Goerg, S. J., & Hamann, H. (2015). Cui Bono, benefit corporation? An experiment inspired by social enterprise legislation in Germany and the US. *Review of Law & Economics*, *11*(1), 811–846.

Glaeser, E., Laibson, D., Scheinkman, J., & Soutter, C. (2000). Measuring trust. *Quarterly Journal of Economics*, *115*(3), 811–846.

Haile, D., Sadrieh, A., & Verbon, H. A. A. (2008). Cross-racial envy and underinvestment in South African partnerships. *Cambridge Journal of Economics*, *32*(5), 703–724.

Hamman, J. R., Loewenstein, G., & Weber, R. A. (2010). Self-interest through delegation: An additional rationale for the owner-agent relationship. *American Economic Review*, *100*(4), 1826–1846.

Hill, A. (2015). Does delegation undermine accountability? Experimental evidence on the relationship between blame shifting and control. *Journal of Empirical Legal Studies*, *12*(2), 311–339.

Johnson, N. D., & Mislin, A. A. (2011). Trust games: A meta-analysis. *Journal of Economic Psychology*, *32*(5), 865–889.

Kvaløy, O., & Luzuriaga, M. (2014). Playing the trust game with other people's money. *Experimental Economics*, *17*(4), 615–630.

Lerdorf, R. (1995). Announce: Personal Home Page Tools (PHP Tools). Current version 5.3 Retrieved from http://www.php.net

Maximiano, S., Sloof, R., & Sonnemans, J. (2013). Gift exchange and the separation of ownership and control. *Games and Economic Behavior*, *77*(1), 41–60.

Mitchell, L. E. (1993). Fairness and trust in corporate law. *Duke Law Journal*, *43*(3), 425–491.

Oexl, R., & Grossman, Z. J. (2013). Shifting the blame to a powerless intermediary. *Experimental Economics*, *16*, 306–312.

Ouma, G. (2014). *University of Pretoria statistics at a glance.* Retrieved from http://www.up.ac.za/media/shared/371/ZP_Files/dipbook_2015.zp43519.pdf. Accessed on May 21, 2018).

Rudman, L. A., & Goodwin, S. A. (2004). Gender differences in automatic in-group bias: Why do women like women more than men like men? *Journal of Personality and Social Psychology*, *87*(4), 494.

Seekings, J. (2008). The continuing salience of race: Discrimination and diversity in South Africa. *Journal of Contemporary African Studies*, *26*(1), 1–25.

Selten, R. (1967). Die Strategiemethode zur Erforschung des eingeschränkt rationalen Verhaltens im Rahmen eines Oligopolexperiments. In H. Sauermann (Ed.), *Beiträge zur experimentellen Wirtschaftsforschung* (pp. 136–168). Tübingen: Mohr.

Tajfel, H. (1970). Experiments in intergroup discrimination. *Scientific American*, *223*, 96–102.

Tajfel, H., & Turner, J. C. (1986). The social identity theory of intergroup behaviour. In S. Worchel & W. G. Austin (Eds.), *Psychology of Intergroup Relations* (2nd ed., pp. 7–24). Chicago, IL: Nelson-Hall.

Van der Merwe, W. G., & Burns, J. (2008). What's in a name? Racial identity and altruism in post-Apartheid South Africa. *South African Journal of Economics*, *76*(2), 266–275.

Zerfu, D., Zikhali, P., & Kabenga, I. (2009). Does ethnicity matter for trust? Evidence from Africa. *Journal of African Economies*, *18*(1), 153–175.

Zizzo, D. J. (2010). Experimenter demand effects in economic experiments. *Experimental Economics*, *13*(1), 75–98.

APPENDIX

COMPUTER PROGRAM MATCHING ALGORITHM

The computer program randomly assigned players into matching groups of six players, where each group included three black and three white players. Three players of each race was the minimum needed to ensure that all players in the group would deal with actual people in all roles (for example, a black player in the role of owner (1) faces a black receiver (2) and chooses between a black sender (3) and a white sender.

When numbers of participants in the session were not an exact multiple of six, a few players were assigned to two matching groups. Since participants were unaware of the existence of the matching groups, this double assignment did not interfere with the running of the game, or with participants' perceptions of the game. The decisions that these players made were simply applied across two groups.

The matching groups were used to randomly assign the owners, senders, and receivers each player would face in the different roles they played. The (behind-the-scenes) matching set-up allowed the program to match amounts sent and returned within each game so that players would be paid according to actual decisions when all payments were revealed at the end of the game. It ensured that players faced either an own- or other-race owner (in the role of receiver) or receiver (in the role of owner); and that they could select from an own and other sender (in the role of owner), where all players were actual individual people in the room.

DEMOGRAPHICS BY RACE GROUP

Table A1. shows the participant demographics for black and white respondents separately.

	Black (%) ($n = 65$)	White (%) ($n = 61$)
Language		
English	25	52
Afrikaans	0	44
Tswana	14	
Zulu	12	
Southern Sotho	11	
Xhosa	8	
Tsonga	8	
Northern Sotho	8	
Other (includes Venda, Swati, Ndebele and 'other')	15	3
Mean age	19.4	19.1
Gender: % female	65	41

ADDITIONAL REGRESSIONS

As per the advice of an anonymous referee, we have included some additional regressions as further robustness checks for our findings.[24] We consider first the standard trust game transfer decision (Table A2); then, the sender decision in the delegation game (Table A3); and finally the receiver decision in the delegation game (Table A4). In each case, we first consider standard group identity (race) variables; we then add expectations about the other player (sender/receiver's decision). In all cases, this significantly improves the model R^2. Finally, we ask whether group identity variables impact expectations. As can be seen in Tables A3 and A4 below, race variables are significant predictors of decisions, in the delegation game, suggesting a role for own-group bias in decision-making. Once expectations of the other player's decisions are included, some of these race-based variables lose significance. However, the race-based variables again enter as significant predictors when we consider predictors of the expectations variable (although the R^2 in these regressions are very low), suggesting that perhaps race-based variables impact behavior through the race-based expectations of the other player's decisions.

Table A2. Standard Trust Game: Sender Transfer Decision.

	Race	Including Expected Return	Predicting Expected Return
Sender race	0.076	0.307	-0.689
	(0.383)	(0.226)	(0.892)
Receiver race	0.127	0.002	0.382
	(0.153)	(0.099)	(0.379)
$S = R$	0.242	−0.022	0.798**
	(0.153)	(0.098)	(0.379)
Expected return		0.329***	
		(0.020)	
session12		−0.152	
		(0.234)	
Constant	5.301***	2.717***	7.985***
	(0.281)	(0.239)	(0.708)
N	504	504	504
adj. R^2	-0.003	0.605	0.002

Notes: Standard errors in parentheses
*$p < 0.10$, **$p < 0.05$, ***$p < 0.01$.

In Table A2, we first regress the transfer amount on the sender race, the receiver race, and whether these align ($S = R$) or not. We then add expected returns to the predictor variables. Finally, we regress expected return on the same predictor variables. These findings show that, while alignment between

sender and receiver race has some impact on expected returns, this does not have a significant impact on transfer decisions.

Table A3. Delegation Game: Sender Transfer Decision.

	Race	Including Expected Return	Predicting Expected Return
Sender race	0.137	0.211	−0.040
	(0.340)	(0.263)	(0.791)
Receiver race	0.281*	0.082	0.520**
	(0.147)	(0.111)	(0.250)
Owner race	0.217*	0.115	0.452**
	(0.130)	(0.100)	(0.224)
$S = R$	0.569***	0.432***	0.527**
	(0.151)	(0.121)	(0.240)
$O = R$	0.240	0.095	1.199**
	(0.166)	(0.136)	(0.513)
session12	-0.288	−0.216	
	(0.375)	(0.291)	
Expected return		0.276***	
		(0.021)	
Constant	5.667***	3.731***	7.147***
	(0.300)	(0.324)	(0.635)
N	1008	955	955
adj. R^2	0.015	0.379	0.010

Notes: Standard errors in parentheses
*$p < 0.10$, **$p < 0.05$, ***$p < 0.01$.

In Table A3, we first regress the sender's transfer decision (lab dollar amount transferred) on the sender's race, the receiver's race, the owner's race, and whether or not the sender and receiver share the same race ($S = R$) and whether or not the owner and receiver share the same race ($O = R$). We then include expected returns as an additional predictor. Finally, we regress the expected return on the same race variables. The regression findings suggest that races of other players, and particularly the parity between the sender and the receiver race has some impact on amounts transferred in the delegation game, likely through the mechanism of differing expected returns from different receiver-owner or receiver-sender pairs.

Table A4. Delegation Game Receiver Return Amounts.

	Race	Including Expected Transfer	Predicting Expected Transfer
Receiver race	−0.371	−0.582	0.154
	(1.011)	(0.815)	(0.380)
Sender race	0.744*	−0.150	0.643**

Table A4. (*Continued*)

	Race	Including Expected Transfer	Predicting Expected Transfer
	(0.408)	(0.251)	(0.276)
Owner race	0.901	0.441	0.331
	(0.719)	(0.536)	(0.311)
$S = R$	0.281	0.090	0.137
	(0.423)	(0.241)	(0.280)
$O = R$	0.176	0.204	−0.020
	(0.406)	(0.262)	(0.243)
session12	0.320	0.252	
	(1.270)	(0.941)	
Expected transfer		1.390***	
		(0.128)	
Constant	6.346***	−1.272*	5.494***
	(0.943)	(0.704)	(0.401)
N	504	504	504
adj. R^2	−0.000	0.385	0.007

Notes: Standard errors in parentheses
*$p < 0.10$, **$p < 0.05$, ***$p < 0.01$.

In Table A4, we first regress the return amount on the races of all players and on whether or not sender and receiver; and owner and receiver races align. We then include the expected transfer as an additional predictor variable. Finally, we regress the expected transfer on the same race variables. Although the sender race is significant in predicting receiver return amounts, it should be noted that this first model has no predictive power. The adjusted R^2 improves once expected transfers are included, and the sender race again appears to play a role in expected transfers.

INSTRUCTIONS TO PARTICIPANTS

Participants received the initial instructions (before Part 1) at the beginning. Following this, the instructions for each part were handed out (and readout) at the beginning of that part. Where different instructions were given for the blind and knowledge treatments, these are highlighted in the below.

Thank you for participating in this decision-making experiment. The following instructions should answer most questions you might have. Otherwise, please raise your hand so that one of the researchers can assist you.

Now that the experiment has begun, we ask the following:

- Please do not talk (except to ask questions of the researchers).
- Please respect the privacy of other participants and do not look at their computer screens.

- Please do not use your cell phone for anything.
- Please do not access any website other than the one you have been directed to.

Should you break any of these rules, you may be excluded from the experiment and any payments.

The experiment consists of a number of parts in which you will have to make decisions about sending money to someone in this room. You will be given limited demographic information about the person concerned in each case. Please note that this demographic information is not detailed enough for you to be identified. Your decisions will therefore be anonymous. Not even the researchers can link your name with the data you provide.

In the experiment you will earn money. How much you earn depends both on your decisions and on the decisions by other participants. Note that, throughout the experiment, you will see all money amounts without a currency. Those amounts are in "lab dollars." Once the experiment is over, the total amount of lab dollars you have earned will be divided by two and paid to you cash in Rands.

Detailed instructions for each part in the experiment follow. Some instructions are included on-screen. The below should offer additional clarity, but if you do have questions not covered by these instructions, feel free to ask the researchers.

Part 1:

In this part you will be paired with another person in this room. You will interact with this person by taking decisions on the computer screen, then your interaction ends. After this, you will be paired with another person and the process repeats. You are free to take different decisions each time and to use the information provided on the computer screen to help you with your decision-making. Part 1 consists of five such interactions, each with a different person in this room.

The first interaction is a test run. It is exactly like the other four, except the "other person" will be simulated by a computerized random algorithm and the round will not be paid. It is merely meant to familiarise you with the basic set-up and to allow for any questions not covered in the text. The first round will be clearly labeled "Test Run" and you will be told the results right afterwards.

The other four interactions are set-up as follows:

- As "Player A," you are given 10 lab dollars and can send any part of this money to Player B.
- The researchers add another two times the amount you sent, so that Player B will receive three times the amount you sent.
- As "Player B," you can send any of the money you received back to Player A.

You will first play the role of Player A, then the role of Player B. After the end of this interaction, you will not be told immediately what the other player

did, but will continue directly with the next interaction. This means that where you are Player A, Player B will not know the amount you sent until the end of the entire experiment. Similarly, when you are Player B, Player A will not know the amount you returned until the very end of the experiment.

At certain points of the interaction, you will be asked what you assume the other player will do. In these cases, please state your honest opinion. Assumptions will never be disclosed to other participants and will only be used by the researchers to better understand your decision-making process.

Part 2: [Number assessment task: not used in this study]

Part 3:
In this part, you will again be paired with one other person. But instead of Player A sending and Player B returning money directly, you will select an extra person in this room (Player X) to take this decision for you. The program will give you a choice between two people in this room, from which you have to select one. This Player X will act on behalf of Player A to make the decision of how much of Player A's money to send to Player B.

So the game now proceeds as follows:

- Player A has 10, some of which can be sent to Player B (as before).
- Player A cannot send money directly, so needs to select a Player X.
- The computer presents two people (candidates) from the room. Player A selects one of these to be Player X. Both candidates for the role of Player X receive 10 lab dollars to keep, no matter which of them is selected to make the decision.
- The selected Player X decides on how much of Player A's money to send to Player B. Player X can choose any amount from 0 to 10. Player X cannot communicate with Player A.
- As before, researchers add money so that Player B receives three times what Player X sent.
- Any money that Player B sends back will go to Player A, as before (NOT to Player X).

You will take decisions in each of the three roles (A, X, B), then the interaction ends and you are paired with different people. There will be four interactions in total (as before, but without the test run) and in each interaction, you will play each of the three roles (A, X, B) once.

Note that where you are Player A, you will have demographic information about the Player B you will be paired with. You will then be able to choose a Player X based on demographic information about two possible Player X options.

[*Know condition instructions:*] Where you are Player B, you will know demographic details about both the Player A who selected the Player X to send you

money; and about the Player X who made the decision about the amount of money to send to you.

[*Blind condition instructions:*] Where you are Player B, you will know demographic details about the Player X who made the decision about the amount of money to send to you. You will not, however, know anything about the Player A who selected the Player X to send you money (remember that the Player A is the person to whom you will be returning money).

Where you are Player X, you will know demographic details about the Player A whose money you are sending AND the Player B to whom you are sending the money (and who will return money to Player A).

SCREENSHOTS SHOWING DECISIONS IN EACH ROLE IN THE DELEGATION GAME

Screenshot showing elicitation of race group information.

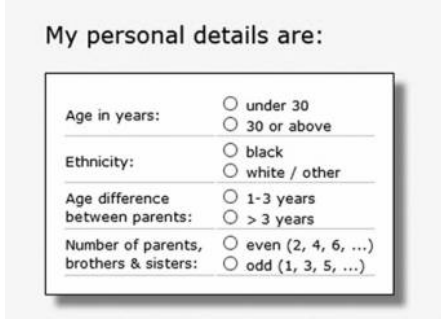

Owner decision (taken from the blind condition, as evidenced by the red text):

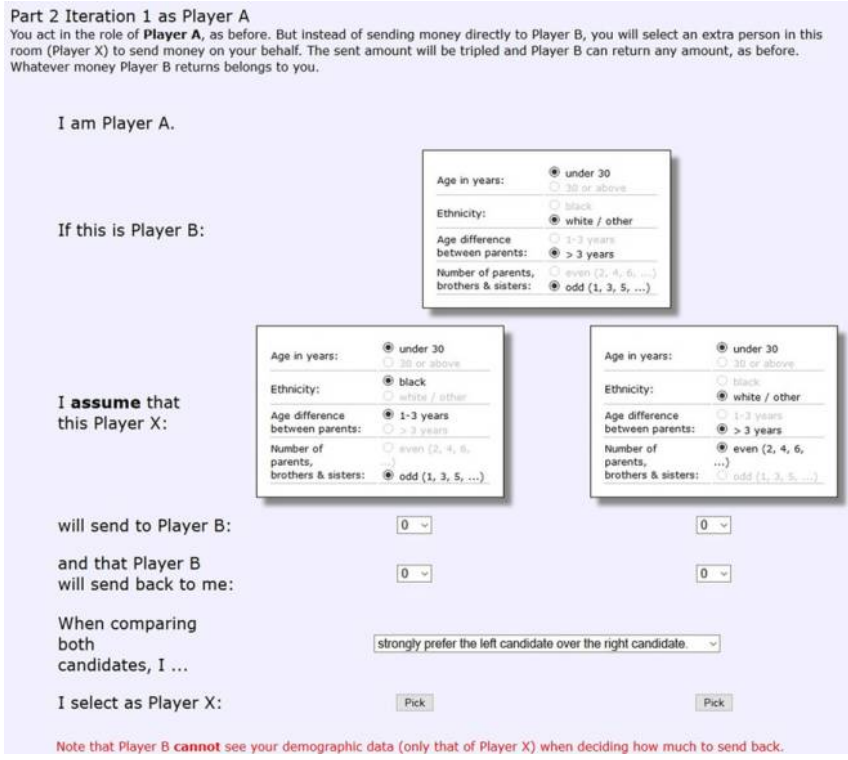

Sender decision (taken from the blind condition, as evident from the red text):

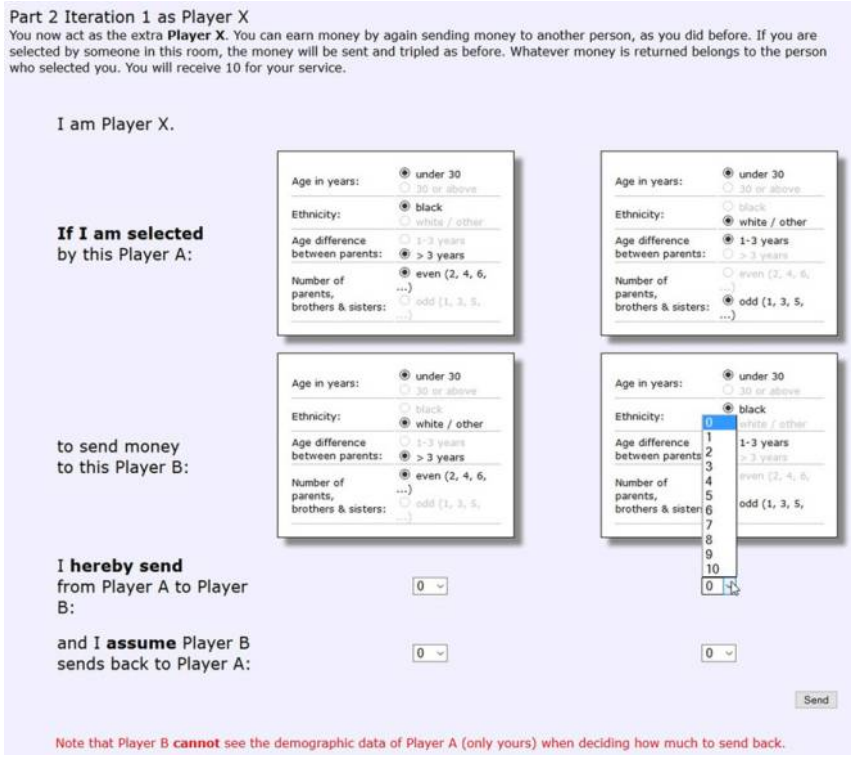

Receiver decision (taken from the blind condition, as evidenced by the red text):

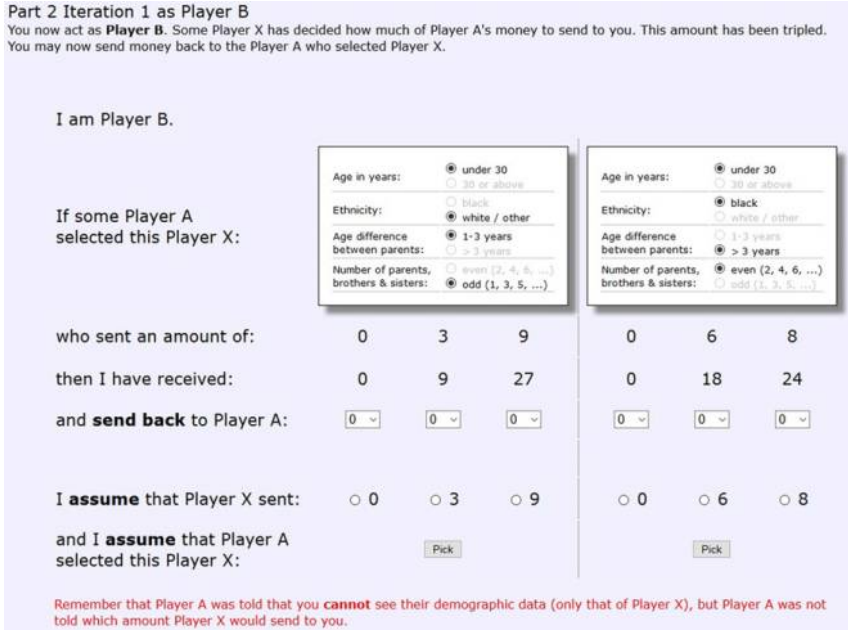

Illustrative extract from demographic questionnaire (at the end of the experiment):

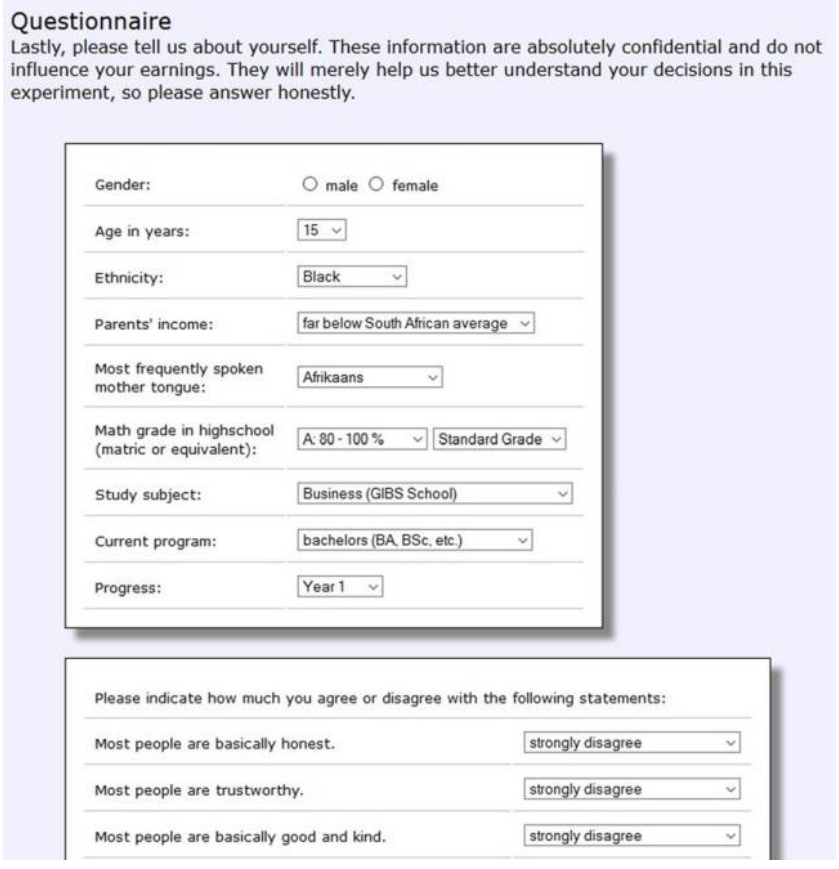

Questionnaire

Lastly, please tell us about yourself. These information are absolutely confidential and do not influence your earnings. They will merely help us better understand your decisions in this experiment, so please answer honestly.

Gender:	○ male ○ female
Age in years:	15
Ethnicity:	Black
Parents' income:	far below South African average
Most frequently spoken mother tongue:	Afrikaans
Math grade in highschool (matric or equivalent):	A: 80 - 100 % / Standard Grade
Study subject:	Business (GIBS School)
Current program:	bachelors (BA, BSc, etc.)
Progress:	Year 1

Please indicate how much you agree or disagree with the following statements:

Most people are basically honest.	strongly disagree
Most people are trustworthy.	strongly disagree
Most people are basically good and kind.	strongly disagree