The Impact of Race and Effort on Social Preferences in an Experimental Setting

UC Berkeley Economics Department: Undergraduate Honors Thesis

Allison Green
Advisor: Stefano DellaVigna

May 5, 2017

Abstract

This paper investigates whether effort, racial identity, and their interaction affect the level of altruism in a dictator game. To test this, I include a randomization step in a dictator game where I vary two characteristics of their receiving partners: their score on a simple data entry task and the race of their partner. The results replicate findings of previous studies to find that as effort increases the overall level of altruism also increases. I also find that black partners on average receive less generosity than their white counterparts. This paper expands upon existing literature by exploring whether the interaction of these two treatments produces additional marginal effects. While in aggregate, there is no marginal effect, when the population is broken down into subgroups of gender, political ideology, and geography, interaction treatment effects are revealed. [1]

[1] I want to sincerely thank Professor DellaVigna for all of his guidance and generous support throughout this thesis process. His help throughout the research design and IRB approval was invaluable for me to complete this thesis.

Contents

1 Introduction and Literature Review

In recent years, developments in behavioral economics established that in social-behavioral setting, people depart from the classical rational theory of economic behavior (Fehr and Schmidt 1999). In particular, one of the most consistently violated axioms of classical theory is selfishness - that an individual only considers their own wellbeing in decision-making. However, while these results are replicated across studies, there is still limited understanding of the determinants of variation in economic altruism.

Existing literature finds that individuals show more altruism towards people within their own social group when those identities are induced in an experimental setting (Chen 2009). However, there is mixed evidence whether natural identities such as race and ethnicity affects social preferences. Glaeser and Soutter (2000) find that in trust games participants give back less money to people of difference races. Bouckaert and Dhaene (2004) use a similar methodology but find giving is independent of ethnic identity. Bernhard and Fehr (2006) show that among different ethnic groups in Papua New Guinea a third-party dictator will be more generous towards people in their own ethnic group. However, a recent study shows that for the majority of the population, intergroup bias is small, but for a distinct subset the bias is extreme enough for participants to sacrifice their own income (Kranton and Huettel 2016). Together this literature suggests that racial or ethnic identity might affect giving for part of the population but not for others. While in-group bias might explain a significant portion of racial discrimination in giving, other seemingly unrelated characteristics might also explain some of the heterogeneity. Additionally, these experiments were all run in different countries with vastly different cultural concepts of race and ethnicity. In the American context, List (2004) in an experiment studying sportscard auctions, finds that individuals belonging to ethnic minorities receive much small offers. Doleac and Stein (2013) find a similar result that black sellers on craigslist receive much lower offers.

There is also some evidence about the effects of effort on experimental altruism. Almås et al. (2007) find that effort significantly motivates giving for part of the population. They divide the population into three subgroups: strict egalitarians, libertarians, liberal egalitarians. The first group are not at all motivated by effort or ability, and conditional on self-interest considerations, they give equally to all people. Libertarians are motivated by effort and ability, and will give higher payouts to partners who produce more. Finally liberal egalitarians, only hold their partners responsible for effort but not ability. They find that a significant portion of their sample population fall into the latter two categories and accordingly modify their level of redistribution based of the level of effort of their partner.

I am interested in not just testing the effects of race/ethnicity and effort on giving separately, but also the interaction of these two treatments. There is no preexisting paper that explores this exact question, so my experiment provides novel insight into the ways these two factors could amplify or mitigate the effects of the other. Previous experimental research shows that individuals have implicit biases towards different racial identities (Doleac and Stein 2013; Bertrand and Mullainathan 2004). These biases are not strictly negative or positive, but act through different channels. One channel biases could work through is the perception of hard work or lack thereof. Specifically, there is a strong unfounded historical-social bias in American history that black individuals have a worse work ethnic than those in other races Katz and Hass 1988). For this experiment, I am exploring whether

this historical-social bias has effects on altruism in an experimental setting.

Beyond the contribution to the field of behavioral economics, these questions of the perception of race and work ethic in the United States have important labor and political economy consequences. In labor settings, if black workers are perceived to be less hard-working or mistakes by them are treated more punitively by superiors, this could be a partial explanation of the wage and promotion discrimination faced by black workers (Grodsky and Pager 2001). The political economy consequences are that preference for redistribution in an experimental setting has an analogous translation to preference for societal redistribution through government intervention. Since black Americans have lower median incomes than average (Census 2010), these questions could provide insight into why certain Americans prefer greater or less levels of redistribution through taxes and government programs.

I set up a behavioral experiment of 818 total participants on Amazon's Mechanical Turk platform which allows for people to complete short tasks for small amounts of compensation. The experiment modified a traditional dictator game by adding a production stage ahead of redistribution where both the dictator and recipient complete a short data task to earn piece rate compensation. The total income from both the dictator and recipient's performance were then pooled together and given to the dictator for redistribution. The dictator had complete control over the final level of redistribution. Dictator games are ideal for measuring l altruism or generosity because dictators have no incentives besides inequality aversion to give their partner anything. I've added a preceding production step as a way of measuring and signaling effort since. This is similar in design to Almås et al. (2007). The data entry task I chose was for respondents to complete repeated CAPTCHA completion (see Gilchrist and Malhotra 2016). I ran the experiment in two stages. The first stage had a total of 236 participants who became the recipient pool. From this pool I generated four effort-race treatment groups. In stage 2, 582 respondents were randomly assigned into different effort-race treatment groups where they each became dictators and chose a level of income redistribution.

In this paper, I will first explore for the pooled-population of dictators the effect of race, effort, and their interaction on altruistic giving. I find that on average effects are small, but this average is not representative. Then, in line with the idea of population heterogeneity for social preferences, I investigate a few different dimensions that could explain some of the heterogeneity. The main dimensions of dictator heterogeneity I am exploring are gender, political ideology, and geographic location. I find the largest heterogeneity in comparing the social preferences of men against women. While male dictators are more generous to white partners than black partners, female dictators' level of generosity is independent of race. I also explore the determinants of a dictator being of the "selfish" type - that they give zero percent of the group income to their partner. While, for the rest of the dictator population, effort is the strongest motivator for increased altruism, the relative level of recipient-effort has no effect on the likelihood of a dictator choosing to be completely selfish.

In section 2 I discuss the models used to formulate the experiment, and the main hypotheses coming from the models. In section 3, I provide a description of the Mechanical Turk platform and the experimental design used. In section 4, I provide analysis of the results, and in section 5 I discuss concluding remarks and talk about potential methodological improvements.

2 Model and Hypothesis

The simplest version of personal utility taking into account fairness considerations for a two-person system is as follows:

$$U_x(x, y) = U_x(x) - \alpha_x Max(0, x - y) - \beta_x Max(0, y - x) \tag{1}$$

adapted from Fehr and Schmidt (1999)

In this model, x is person x's income and y is the person y's income. $U_x(x)$ is the utility of person x independent of the income of person y, and $U_x(x, y)$ is the utility of person x incorporating the income of person y. α_x measures the level of economic altruism towards person y. We assume that it is greater than or equal to 0 and less than 1. Intuitively, β_x represents the level of redistribution resulting from a desire for fairness person x feels when they are better off than person y. is person x's inequality aversion to person y. We assume that it is greater than or equal to 0 and less than 1. Intuitively, it represents person x's desire for redistribution due to fairness when person y is better off than person x.

Across many lab settings, it has been established through behavioral games such as a Dictator game, that α_x is greater than 0 (Camerer 2003). There is also a breadth of literature showing that α_x varies both with the characteristics of person x and the characteristics of person y (examples include Bouckaert and Dhaene 2004; Almås et al. 2007; Kranton and Huettel 2016).

For the purpose of this paper, I am interested in two characteristics about person y that could affect α_x : effort and race/ethnicity. Consider the following model to determine α_x for person x in a simplified setting where the only other person is person y.

$$\alpha_x = f(E_y, R_y, O_y, X) \tag{2}$$

E_y is the effort level of person y as perceived by person x, and R_y is the race of person y as perceived by person x. O_y is the set of all characteristics of person y that are observed by person x besides race and effort. Depending on the setting, these could include characteristics such as age, gender, attractiveness, likability, or intensity. Importantly, these variables are all person x's perception of person y, and they could be different from the true value of person y. X is all of the characteristics of person x observable and unobservable that affect their level of altruism.

I make the following assumptions about equation 2:

- The covariants are not likely to be independent from each other.
- The function can be approximately modeled with linear terms and linear interaction terms.

From these assumptions I deduce the following model:

$$E(\alpha_x | O_y; X) = \beta_0 + \beta_1 E_y + \beta_2 R_y + \beta_3 (E_y * R_y) \tag{3}$$

Since the experiment is being run anonymously over an online platform, I can closely control for the information about person y that person x is perceiving. Since the only two

pieces of information about person y that is being supplied to person are the photograph of the hand and their score on the CAPTCHA task any variables in O_y would be coming from that information. Therefore, we make the assumption:

$$O_y = \emptyset$$

Additionally, I randomized the dictators into different treatments, so we can make the following assumption:

$$E(X)_j = E(X)_i$$

For treatments i and j. Using all of these assumptions I deduce the following empirical model for testing the true value of α_x for person x.

$$E(\alpha_x|O_y; X) = E(\alpha_x) = \beta_0 + \beta_1 E_y + \beta_2 R_y + \beta_3 (E_y * R_y) \tag{4}$$

In the case of this experiment, I have simplified E_y and R_y into binary variables. E_y takes on 1 when person y is in the top 20% of scores on the CAPTCHA task, and 0 when person y is in the bottom 20% of scores on the CAPTCHA task. Respondents in the middle 60% of the distribution were not included in the experiment. R_y takes on the value of 1 when person y is identified as white and 0 if identified as black. Respondents not identifying as either were excluded in the experimental phase. α_x is measured by the percent of total income redistributed from person x to person y. β_3 can be interpreted as the marginal expected payout of being both white and top-scoring in the CAPTCHA task.

2.1 Main Hypotheses

- **Hypothesis 1**: When $E_y = 1, \alpha_x$ will be higher.
  When the effort level of person y is high, the amount of income distributed to them will be higher. This result would be consistent with existing literature. (see Almås et al. 2007)

- **Hypothesis 2:** When the race of person y is black, α_x will be lower.
  This is supported by two results in current literature. First, literature on behavioral racial discrimination (Doleac and Stein 2013; Bertrand and Mullainathan 2004) would suggest that the whole population would be less generous to partners they perceive as black. Secondly, literature on group behavior (Kranton and Huettel 2016; Fershtman and Gneezy 2001; Bouckaert and Dhaene 2004 would suggest a population of dictators that is majority white (as is the case with Mechanical Turk) would be more generous towards other white participants.

- **Hypothesis 3** The interaction term β_3 will be positive. The interaction term measures the size and direction of that effect. When β_3 is positive, as effort on the CAPTCHA task increases, payouts to white partners increase relative to black partners.

- *Interpretation 1:* In this case, dictators are highly sensitive to the productivity of white partners, but relatively insensitive to the relative productivity of black partners. This scenario occurs if dictators have higher expectations for the performance of white partners than they do black partners. When white partners underperform to these standards dictators have more additional negative feelings towards low effort white partners relative to high effort white partners than they do for black partners. Evidence for this case would be if low effort white partners were paid similarly or even worse than low effort black partners, and black partners were paid similarly regardless if they were high effort or low effort.
- *Interpretation 2:* Dictators are sensitive to the productivity of black partners, but actually pay black partners worse when they do better on the CAPTCHA task. In this situation, dictators feel negatively about black partners outperforming them on a task, and respond with punitive behavior. Evidence for this situation would be if payouts to white partners stay the same or increase as effort increases, but payouts to black partners actual decrease as effort increases.

- **Hypothesis 4:** The interaction term β_3 will be negative. This is the opposite effect as anticipated by hypothesis 3. Therefore, as the overall effort level of a paired-partner increases, that payouts to white partners decrease relative to black partners.
  - *Interpretation 1:* Dictators are sensitive to the effort of black partners, and relatively insensitive to the productivity of white partners. Evidence for this interpretation would be payouts to black partners increasing with effort, and payouts to white partners staying relatively stagnant or even decreasing. The channel of bias this could be working through is that high-performance among black partners is more surprising, and therefore it triggers a stronger response from the dictator in terms of their fairness considerations.
  - *Interpretation 2:* Dictators are sensitive to the effort of white partners, and relatively insensitive to the productivity of black partners. In this case, payouts to white partners would actually decrease with effort, while payouts to black partners would stay relative stagnant or even slightly increase.

For hypotheses 3 and 4, there is also the possibility that both interpretations could be true and contributing to the overall magnitude and directionality of β_3 . I will also investigate whether for some subgroups of people hypothesis 3 holds true, and for other groups of people hypothesis 4 holds true. This would be consistent with the findings of Kranton and Huettel (2016) that when the population is broken down there are different subgroups that have opposite reactions to treatment.

3 Experimental Design

The experiment was designed to modify a traditional dictator game to investigate two potential determinants of payouts in a dictator game: effort and race. In a traditional two-person dictator game, one person is randomly given a lump sum of income and given the power to choose a final level of income distribution between both people. The major modification for this experiment are adding a data entry task with piece rate compensation before (see Almås et al. 2007), and pooling both player's income from the task to create the income to be redistributed. The other modification is showing a photograph of the recipient's hand to the dictator before they choose redistribution to signal racial identity (see Doleac and Stein 2013). The experiment was run in two phases where the first phase was used to generate the recipient pool, and the second phase was the experimental phase where the dictators made their payoff decision. The decision to run the experiment in two stages with two separate groups of respondents was mostly due to the the disproportionately small number of black individuals using the Mechanical Turk site.

3.1 Mechanical Turk Platform and Subject Pool

The experiment was run on Amazon's Mechanical Turk platform (MTurk). This platform allows anyone to upload small microtasks in batches for people in the MTurk worker pool to complete. This platform is becoming increasingly popular in economic research (DellaVigna and Pope 2016) as an alternative to traditional laboratory research pools due to its inexpensive costs.

Goodman and Cheema (2013) evaluate the strengths and weaknesses of using mTurk for social science experiments relative to the traditional college pool. The major drawback they highlighted was worker inattentiveness during tasks. They also noted that mTurk workers are significantly less extroverted and having lower self-esteem than the general population.

Beyond cost efficiency, Mechanical Turk offers a few other benefits compared to traditional laboratory experiments. For one, many laboratory experiments take place at universities with student subject pools that are very homogenous in terms of age, education level and relatively homogenous in terms of family income level and race. The Mechanical Turk sample, while not completely representative of the United States is much more representative than a college pool.

Overall Goodman and Cheema (2013) find that the benefits outweigh the drawbacks, and Mechanical Turk represents a suitable if not preferable platform for behavioral research.

3.2 Data Entry Task - CAPTCHA completion

As a proxy for effort, I used a data entry task that all participants in stages one and two completed. The ideal data entry task would be one where the only determinant of performance on the task is effort. I chose a CAPTCHA completion task t was something that almost all of the subject pool would have experience working with. Gilchrist and Malhotra (2016) use a similar task a proxy for respondent effort. Our methodologies diverge in that in their task participants were asked to type CAPTCHAs for multiple hours, while my task only lasts for 2 minutes. Therefore conditional on effort, typing speed (a preexisting skill) could

affect performance and more importantly the perception of performance. See a discussion of possible methodology flaws and improvements in the conclusion.

Each subject was asked to type CAPTCHAs for two minutes, and offered a piece-rate compensation $0.01 per CAPTCHA that they accurately typed. All of the CAPTCHAs were organized on one page, and each CAPTCHA consisted of two randomly generated words (see Appendix 1 for example). The CAPTCHAs were the same set of 60 CAPTCHAs ordered in the same way to maximize comparability across observations. Using a scoring mechanism on Qualtrics, each participant was given a score of between 0 and 60.

Since participants in stage one were ineligible to participate in the second stage, there is no concern about someone having the opportunity to improve upon their previous CAPTCHA score by typing the same set of words twice.

3.3 Recipient Generation Stage and Treatment Creation

I recruited 236 participants in the first stage to generate the partners for matching in stage 2. In this stage all participants first completed the CAPTCHA task, and then they were asked to upload a photograph of their hand holding the number of CAPTCHAs they completed. Doleac and Stein (2013) use photographs of hands as a proxy for race and find that it effectively conveys racial identity to participants. Following this, participants were asked to answer a few demographic questions on race,age,and gender. Full completion of the survey took on average 8 minutes.

Based on their answer to on their score, each of the participants was categorized into one of five groups: high effort and black (Treatment 1), high effort and white (Treatment 2), low effort and black (Treatment 3) , low effort and white (Treatment 4), and everyone else (no treatment). A score was identified as high effort if it was in the top 20 percent of scores and low effort if in the bottom 20 percent. In the recipient pool, the 20th percentile score was 21, and the 80th percentile score was 39. Assignment to black and white was based on the self-identification of the participant. All stage 1 respondents not sorted into treatments 1-4 were dropped from being used in stage 2

The people in these four treatments were matched with partners in the second stage. Because of the relatively larger number of subjects in the second stage, each person was matched with multiple people in stage 2 and they were randomly allocated one payout of all the payouts given to them. Because there is unequal amount of subjects in each treatment, certain subjects were assigned to each treatment in the second phase.

3.4 Experimental Stage

In the experimental stage, the survey was open to the same population as in the recipient generation stage except people who participated in the first stage were not eligible for the second stage. Each person participating in this stage automatically assumes the role of distributor/dictator in the dictator game. 582 participants were recruited into this. They were first asked to do the CAPTCHA completion, then based on random assignment explained above, they were assigned a partner. Full completion of the survey took on average 7 minutes.

Based on their treatment assignment, they were given three pieces of information: their partner's score from the CAPTCHA task (to signal effort) , the total combined income, and the photograph of their partner's hand (to signal race). Appendix 2 contains the survey page the dictators/distributors were presented with for one of the treatments. Doleac and Stein (2013) similarly use photographs of hands in craigslist adds to signal race in an online experiment. I set the counter on the survey to keep all respondents on this page of the survey for 15 seconds to reduce the likelihood that they would skip to the next page without absorbing the intervention information.

On the next page of the survey, dictators were asked to choose what percentage of the total income to give to their partner. The slider the participants used to decide was automatically to 50 percent. To reduce the biasing effect of this, I set the survey so dictators/distributors could not progress to the next page for 20 seconds, and made it so they had to click on the counter for it to register a distribution amount.

4 Results and Discussion

4.1 Descriptive Statistics

Table 1 presents the frequency and proportion of different payouts from the experimental stage in intervals of 10 and Figure 1 presents a histogram of payouts. The mean percent given to partners was is 30.6%, and the median is 32.5%. The minimum given was 0% and the maximum was 100%. This is higher than results from general dictator games, but it is consistent with the level of giving found from other dictator games with a preceding production phase (Almås et al. 2007).The most common payouts were 0% and 50%, and the distribution of payouts show that dictators/distributor have a strong preference for round numbers.

32% of all dictators/distributors gave their partners nothing. 17% of dictators/distributors gave their partners 50% of total earnings. I have categorized dictators/distributors into 4 broad social preference categories: selfish, somewhat altruistic, egalitarian, and altruistic. Selfish dictators (30% of sample) give zero percent of pooled earnings back to their partner. Altruistic dictators (21%) give more of the pooled income to their partners than they keep for themselves. Somewhat altruistic dictators (32%) keep more of the pooled income for themselves, but they give some of it to their partners. Egalitarian dictators (17%) give exactly half of pooled income to their partners. Figures 3 and 3a show the distribution of the different types of dictators by treatment and gender. Figure 3 shows that dictators in treatment 1 and 2 (the high effort treatments) are roughly twice as likely to be altruists. However, dictators in treatment 1 and 2 are also slightly more likely to be selfish. People in treatments 3 and 4 (low effort) are more likely to be partial altruists. Figure 3a illuminates some visual differences in the gendered response to the different treatments.

Table 2 presents summary statistics for the population in aggregate and the sub-populations of male and female dictators. Woman on average gave 4 percentage points more in payouts then men, and this difference is statistically significant at the 10% level. On the political spectrum, the average ideology was slightly more than 3 (slightly liberal) indicating that the sample of respondents skew slightly left. This is consistent with other surveys on the

political ideology of the mTurk sample population (Huff and Tingley 2015). The average age of respondents was 36, and the women in sample were on average 4 years older then the men in the sample.

Figure 2 presents a histogram of scores on the CAPTCHA task for dictators/distributors. The average score on the CAPTCHA task for the experimental stage was 26. This is significantly less than the average score for the recipient group which was 30. Table 3 shows a comparison of the quintiles for the dictator group compared to the recipient group. Figure 4 shows a comparison of the density plots for the scores of dictators and recipients on the CAPTCHA task.

For all subsequent analysis, I separated the dictator population into two groups: completely selfish dictators (gave 0 percent) and partially altruistic dictators (gave greater than 0 percent). This was motivated by that these two groups are fundamentally different in motivation, and therefore would react differently to the intervention.

4.2 General Regressions

Table 3 presents the general treatment effects of an OLS model with robust standard errors for the sample of dictators who gave more than 0 percent in the redistribution phase. Overall, these effects are muted. Effort alone increases generosity by 11.6 percentage points providing support for hypothesis 2. Race alone does not significantly increase generosity, but the point estimate is positive indicating that there might exist small racial bias in favor of white partners. When the interaction term is included, the term for effort actually increases. The point estimate for the interaction term is small in magnitude and not statistically significant. This provides no support for hypothesis 3. These results suggest that as a whole population effort motivates increased giving, but racial identity does not.

Table 5 presents the outputs of a fitted logit model for the likelihood of giving 0 percent in the redistribution phase. I used a logit specification because the outcome variable is binary. By this estimation the race and effort treatments had no significant impact on the likelihood of a dictator giving their partner nothing. The point estimates are small, but show that white partners are slightly less likely to be given nothing. Unlike in the general model, the point estimate for effort indicates that if anything being high effort makes the dictator slightly more likely to give zero payout. One possible explanation for this could be that the higher pooled income becomes a more enticing prospect that a lower pooled income.

4.3 Heterogeneity in Gendered Giving

There is existing literature establishing that men and women have different patterns of generosity in experimental setting. Andreoni and Vesterlund (2001) find that men are more likely to be completely selfish or completely selfless, while women are more likely to be egalitarian or somewhat altruistic. My results support that men are more likely to be completely selfish but it does not support that they are more likely to be completely altruistic.

Table 4 shows the fitted OLS models with robust standard errors for men and women who gave more than 0 percent. While men give white partners 5.4 percentage points more in payouts, race has almost no effect on the altruism of women. However, while women give

high-effort partners 13.1 percentage points more, men give 11.5 percent points suggesting that female are more responsive to the effort of their partner.

When the interaction term is included, the sign of β_3 for men and women is opposite. For women, the interaction term is negative providing support for hypothesis 4 - that as effort increases the payouts to black partners relative to white partners increase. Because woman actually pay high-effort black partners more than high-effort white partners this complies with interpretation 1 of hypothesis 4. Overall, the interaction term is statistically significant, so while it shows some evidence for this effect it cannot be asserted with any confidence.

For men, the included interaction term is positive providing some support that hypothesis 3 holds for men. Among both the high effort and low effort pools, men pay white partners better than black partners. However, they pay high effort white partners twice as much more than high effort black partners relative to their low effort counterparts. Overall this provides support for interpretation 1 of hypothesis 3. Together these results suggest that there is definite heterogeneity in the social preferences of men and women in the sample.

Table 6 shows that unlike in the general pooled model, there are statistically significant treatment effects on the log odds of giving nothing depending on the race and effort level of a dictator's paired partner. For men, the point estimates on race and effort indicate that men are more likely to give black partners nothing, and they are also more likely to give high effort partners nothing. Neither of these point estimates are statistically significant. When the interaction term is included it provides additional support that male respondents are following the behavior described in interpretation 1 of hypothesis 3.

For women, the point estimates on race and effort indicate that women are slightly more likely to give white participants nothing and slightly less likely to give high effort participants zero. When the interaction term is included, the coefficients become more strongly statistically significant. Among "selfish"-type women, the interaction term shows strong support for interpretation 1 of hypothesis 4. This means that in the low-effort pool, women are much more likely to give white partners nothing, but in the high effort pool women are slightly more likely to give black partners nothing. The large magnitude of β_3 indicates that women of the completely selfish are much more responsive to white effort than black effort.

All together these results provide further evidence that men and women have different social preferences. Specifically, women are slightly more generous overall and more motivated by the effort level of their partner than their race. Men, on the other hand, are more generous to white partners than black partners. While they also are responsive to the effort level of their partners, the magnitude is smaller than for women.

4.4 Other Heterogeneity: Political Ideology and Geography

I provide initial analyses on the effects of variation of political ideology and geography on altruistic giving.

Political Effects

In a social psychology paper, Graham and Nosek (2009) find that liberal and conservatives rate moral intuitions differently. While those who identify as liberal prioritize fairness/reciprocity more, conservatives prioritize in-group/loyalty more. This could indicate in terms of this experiment that one, liberals will be overall more generous overall, and two

that conservatives might be marginally more prone to in-group bias. In this case since the majority of the dictator pool is white, this could translate to greater generosity to white partners relative to black partners.

For the purpose of specification, I label all respondents who rated themselves as 5 (slightly conservative) or greater as the "conservative" group, and I label all respondents who rated themselves as 3 (slightly liberal) or less as the "liberal" group. The respondents who rated themselves as a 4 (moderate) were removed for this portion of analysis. Table 7 includes the regression results of OLS regressions on the liberal and conservative subgroups. The mean percent given in the liberal group was 31 percent and the mean percent given in the conservative group 29 percent. This is somewhat consistent with my stated hypotheses, but the two groups are pretty similar in general giving.

Overall, the major result from these specifications is that the conservative population is more sensitive to the effort intervention than the liberal population. In the conservative group, being matched with a high effort partner increased giving by 13 percentage points, while in the liberal group being matched with a high effort partner only increases payouts by 8 percentage points.

In terms of racial effects, the point estimates were small in magnitude and not statistically significant. When the interaction term is included, the interaction coefficient is small in magnitude and not statistically significant for groups. Therefore, based on these specifications the results do not show any relationship between a dictator's perception of their's partners race, their political affiliation, and generosity.

Geographic Effects

In terms of geographic effects, I originally intended to do analysis on the county level, but a significant portion of the respondent pool misinterpreted the question and answered it incorrectly. Therefore, I updated my analysis to reflect by just exploring the consequence of variation in states.

Red States vs. Blue States

I am interested by looking at differences between people in red states vs. blue states independent of their own personal leanings. The purpose of this specification is to see if cultural differences across states affect the impact of treatments on generosity and fairness considerations. Ideally, this would be studied at the congressional level, but for a rough analysis I am using a state's voting outcome in the 2016 Presidential election.

Interestingly, the mean political score in the subset of dictators living in red states was 3.3, and the mean for those living in blue states was 3.1. Therefore, the actual political leanings of individuals living in red vs. blue states are pretty similar in aggregate. This is good for analysis because it means the treatment effects I'm capturing are somewhat independent of personal political preferences.

Table 8 shows the results from an OLS regression with robust standard errors on the dictator group separated by red states and blue states. Like in earlier specifications, I excluded the "selfish"-type from this specification and studied them separately. In red states dictators give white partners 7.5 percentage points higher payouts on average than black partners, while in blue states dictators give white partners 3 percentage points lower payouts than black partners on average. This is in contrast to table 7 where was no clear difference between conservative and liberals. One interpretation of this is that implicit biases having to do with race have more to do with geography than political ideology. Another

interpretations is that some people inaccurately self-assess their own political ideology (either through careless survey taking or other issues) , and geography is more accurately at roughly estimating political ideology.

Table 9 shows the log likelihood of being a "selfish" dictator for people living in red states and blue states. There are no major statistically significant results for these specifications, and the magnitude of all the point estimates are relatively small.

5 Conclusion

To study heterogeneity in preference for altruistic giving, I conducted a modified dictator game on the Mechanical Turk platform. By including two different randomized treatments (race and effort) and interacting them, I was able to confirm and expand upon the existing literature of Kranton and Huettel (2016); Almås et al. (2007); Chen (2009), by not just showing that altruistic giving varies with different partner characteristics, but that intersection of these characteristics create additional marginal treatment effects.

Taken together, results suggest that there is some relationship between effort, race, and the interaction of them on the distributive preferences, but the high variance in response data make it impossible to draw any definitive claims on the magnitude or strength of this relationship. While the effect of the effort intervention was consistent in the general population and across subgroups, the impact of the race intervention varied across different groups. By looking at the variation of the effect of the intervention on different subgroups of the sample population, it is clear that there is a high degree of heterogeneity in the way partner effort and race affect a person's level of altruism. In particular, men and women exhibit different responses to the interventions. While men were more sensitive to the race intervention, and tended to be more generous towards white partners, women were more sensitive to the effort intervention, and were most generous to high-scoring partners independent of race. Further analysis using state geography showed that independent of political ideology, dictators living in red states were less generous to black partners relative to their counterparts living in blue states.

While the experiment yielded some promising results, a few methodological improvements in a future experiment could provide clearer insight into understanding these relationships more fully.

One substantial improvement would be to make the effort task longer. A longer effort task would make the skill aspect of the task (typing speed) more relevant because if a person was asked to type captchas nonstop for 10 minutes rather than 2 minutes, this would probably differentiate people more on their desire to push through a tedious task rather than just the speed at which a person can type. This would also make the effort task more in line with the practices of Gilchrist and Malhotra (2016)

Another improvement would be to include questions at the end of the survey to more explicitly measure whether either the effort intervention and race intervention effectively alter the perception of the dictator. I would also suggest investigating other ways of signaling race through an online experiment other than using a hand. By using multiple treatment arms that signal race in different ways, it would be possible to more accurately discern the effect

of racial identity on altruistic giving.

Additionally, the experiment was not explicitly set up to measure in-group favoritism, but the racial demography of Mechanical Turk produced a dictator sample that was majoritarian white. A further experiment expanding upon these findings could more carefully thinking about how to work around these issues.

Since there was such a distinct measured difference of the treatment effect on men and women in the sample, further research and replication could more explicitly develop an answer to whether these differences represent a true difference in the American population. Since with subsetting and excluding the "selfish" type, the samples in the regression were relatively small, it could be that the measured effect was just a quirk of the sample.

6 References

References

Almås, I., A. W. C. E. Ø. Sørensen, and B. Tungodden (2007). The pluralism of fairness ideals: An experimental approach. *The American Economic Review 97*(3), 818–827.

Andreoni, J. and L. Vesterlund (2001). Which is the fair sex? gender differences in altruism. *The Quarterly Journal of Economics 116*(1), 293–312.

Bernhard, Helen, U. F. and E. Fehr (2006). Bernhard, helen, urs fischbacher, and ernst fehr. *Nature 442*(7105), 912–915.

Bertrand, M. and S. Mullainathan (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.

Bouckaert, J. and G. Dhaene (2004). Inter-ethnic trust and reciprocity: results of an experiment with small businessmen. *European Journal of Political Economy 20*(4), 869–886.

Camerer, C. (2003). *Behavioral game theory: Experiments in strategic interaction*. Princeton University Press.

Census, U. (2010). United states population census.

Chen, Y. . S. X. L. (2009). Group identity and social preferences. *The American Economic Review 99*(1), 431–457.

DellaVigna, S. and D. Pope (2016). What motivates effort? evidence and expert forecasts. *National Bureau of Economic Research* (w22193).

Doleac, J. L. and L. C. Stein (2013). The visible hand: Race and online market outcomes. *The Economic Journal 123*(572), F469–F492.

Fehr, E. and K. M. Schmidt (1999). A theory of fairness, competition, and cooperation. *Quarterly journal of Economics*, 817–868.

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

Gilchrist, Duncan S., M. L. and D. Malhotra (2016). When 3+ 1¿ 4: Gift structure and reciprocity in the field. *Management Science*.

Glaeser, Edward L., D. I. L. J. A. S. and C. L. Soutter (2000). Measuring trust. *Quarterly Journal of Economics*, 811–846.

Goodman, Joseph K., C. E. C. and A. Cheema (2013, 224). Data collection in a flat world: The strengths and weaknesses of mechanical turk samples. *Journal of Behavioral Decision Making 26*(3), 213.

Graham, Jesse, J. H. and B. A. Nosek (2009). Liberals and conservatives rely on different sets of moral foundations. *Journal of personality and social psychology 96*(5), 1029.

Grodsky, E. and D. Pager (2001). The structure of disadvantage: Individual and occupational determinants of the black-white wage gap. *American Sociological Review*, 542–567.

Huff, C. and D. Tingley (2015). "who are these people?" evaluating the demographic characteristics and political preferences of mturk survey respondents. *Research & Politics 2*(3).

Katz, I. and R. G. Hass (1988). Racial ambivalence and american value conflict: Correlational and priming studies of dual cognitive structures. *Journal of personality and social psychology 55.6*.

Kranton, Rachel, M. P. S. S. and S. Huettel (2016). Groupy and non-groupy behavior: Deconstructing bias in social preferences. *Discussion Paper, Duke University*.

List, J. (2004). The nature and extent of discrimination in the marketplace: Evidence from the field. *The Quarterly Journal of Economics 119*(1), 49–89.

7 Tables and Figures

Table 1: Frequency Table by Payoff

Percent	Frequency
0	176
10	27
20	51
30	26
40	48
50	124
60	54
70	17
80	13
90	3
100	9

Table 2: Summary Statistics

	Full Sample	Women	Men
Political	3.17	3.07	3.25
Age	36.49	38.65	34.70
Score	26.05	25.13	26.82
Eduction - Proportion by level			
Associate degree in college (2-year)	0.13	0.13	0.13
Bachelor's degree in college (4-year)	0.39	0.40	0.38
Doctoral degree	0.02	0.02	0.02
High school graduate	0.11	0.09	0.12
Master's degree	0.09	0.09	0.08
Professional degree (JD, MD)	0.01	0.02	0.01
Some college but no degree	0.25	0.25	0.26
Income - Proportion by bracket			
$10,000 - $29,999	0.24	0.22	0.25
$100,000 - $149,999	0.06	0.06	0.06
$150,000 - $249,999	0.03	0.02	0.03
$30,000 - $59,999	0.34	0.31	0.36
$60,000 - $89,999	0.17	0.21	0.14
$80,000 - $99,999	0.10	0.10	0.10
Less than $10,000	0.06	0.07	0.06
More than $250,000	0.01	0.01	0.01
Dictator Race/Ethnicity - Proportions			
American Indian or Alaska Native	0.00	0.00	0.00
Asian	0.07	0.06	0.08
Black or African American	0.08	0.10	0.07
Hispanic, Latino, or Spanish Origin	0.04	0.03	0.04
Middle Eastern or North African	0.00	0.00	0.01
Native Hawaiian or Pacific Islander	0.00	0.00	0.00
Other	0.01	0.02	0.00
White	0.76	0.75	0.78
White,American Indian or Alaska Native	0.00	0.00	0.00
White,American Indian or Alaska Native,Hispanic, Latino, or Spanish Origin	0.00	0.00	0.00
White,Asian	0.01	0.02	0.00
White,Black or African American	0.01	0.00	0.01
White,Black or African American,Other	0.00	0.00	0.00
White,Hispanic, Latino, or Spanish Origin	0.01	0.01	0.00

Note: Table 2 presents summary statistics for the full sample of dictators, and the sample of dictators separated into male and female subsets. This table includes summary statistics of the demographic questions asked post-experiment.

Table 3: General Treatment Effects - OLS

	Dependent variable:		
	Race	percent Effort	Interaction included
	(1)	(2)	(3)
white	3.028		2.985
	(2.035)		(2.534)
effort		11.615***	12.027***
		(1.849)	(2.655)
whitexeffort			−0.898
			(3.661)
score	−0.619***	−0.609***	−0.626***
	(0.092)	(0.083)	(0.083)
Observations	383	383	383
Residual Std. Error	17.861 (df = 365)	15.852 (df = 365)	15.028 (df = 363)

*p<0.1; **p<0.05; ***p<0.01

Note: Table 3 shows the result of 3 different OLS model specifications with robust standard errors. For each model the sample includes all dictators who gave more than zero percent in the redistribution phase of the survey. Controls include in this specification are: age, score on CAPTCHA task, education level, income, and self-described political leaning

Table 4: Gendered Treatment Effects - OLS

	Dependent variable:					
	percent					
	Men - Race	Men - Effort	Men -Interaction	Women - Race	Women - Effort	Women -Interaction
	(1)	(2)	(3)	(4)	(5)	(6)
white	5.413**		3.836	−0.324		1.523
	(2.614)		(3.093)	(3.470)		(4.439)
effort		11.468***	9.408***		13.167***	16.821***
		(2.200)	(3.217)		(3.070)	(4.697)
whitexeffort			3.945			−6.419
			(4.519)			(6.284)
Observations	196	196	196	187	187	187
Residual Std. Error	14.209 (df = 178)	12.279 (df = 178)	12.906 (df = 176)	21.037 (df = 169)	17.691 (df = 169)	18.448 (df = 167)

*p<0.1; **p<0.05; ***p<0.01

Table 4 shows the results of 6 different specification for OLS estimation of treatment effects with robust standard errors. For each model the sample includes all dictators who gave more than zero percent in the redistribution phase of the survey. Controls include in this specification are: age, score on CAPTCHA task, education level, income, and self-described political leaning

Table 5: Treatment Effects for "Selfish" Type - Logit

	Dependent variable:		
	zero		
	Race	Effort	Interaction included
	(1)	(2)	(3)
white	−0.050		0.011
	(0.196)		(0.283)
effort		0.224	0.285
		(0.195)	(0.279)
whitexeffort			−0.122
			(0.399)
Observations	548	548	548
Log Likelihood	−321.837	−321.206	−321.126
Akaike Inf. Crit.	705.673	704.412	708.252

*p<0.1; **p<0.05; ***p<0.01

Note: Table 5 shows the results of a logit estimation of the treatment effects for the subset of dictators that gave zero percent in redistribution. Each coefficient measures the log likelihood of giving zero based off the binary being on or off. Controls include in this specification are: age, score on CAPTCHA task, education level, income, and self-described political leaning

Table 6: Gendered Treatment Effects for "Selfish" Type - Logit

	Dependent variable:					
				zero		
	Men - Race	Men - Effort	Men -Interaction	Women - Race	Women - Effort	Women -Interaction
	(1)	(2)	(3)	(4)	(5)	(6)
white	−0.157		−0.538	0.273		1.208**
	(0.263)		(0.381)	(0.327)		(0.504)
effort		0.442*	0.082		−0.054	0.940*
		(0.260)	(0.357)		(0.326)	(0.520)
whitexeffort			0.779			−1.775**
			(0.540)			(0.694)
Observations	300	300	300	248	248	248
Log Likelihood	−181.073	−179.793	−178.589	−126.168	−126.506	−122.728
Akaike Inf. Crit.	416.146	413.585	415.179	308.336	309.011	305.457
Note:						*p<0.1; **p<0.05; ***p<0.01

Note: Table 6 shows the results of a logit estimation by gender of the treatment effects for the subset of dictators that gave zero percent in redistribution. Each coefficient measures the log likelihood of giving zero based off the binary being on or off. Controls include in this specification are: age, score on CAPTCHA task, education level, income, and self-described political leaning

Table 7: Political Effects

	Dependent variable:					
	percent					
	Conservative - Race	Conservative - Effort	Conservative -Interaction	Liberal - Race	Liberal - Effort	Liberal-Interaction
	(1)	(2)	(3)	(4)	(5)	(6)
white	2.205		0.377	1.016		1.004
	(5.162)		(6.735)	(2.540)		(3.451)
effort		13.178**	12.281*		7.666***	7.568**
		(5.029)	(7.141)		(2.420)	(3.468)
whitexeffort			1.599			0.222
			(9.950)			(4.922)
Observations	99	99	99	239	239	239
R^2	0.333	0.388	0.388	0.197	0.233	0.233
Adjusted R^2	0.129	0.200	0.179	0.094	0.134	0.127
Residual Std. Error	22.771 (df = 75)	21.821 (df = 75)	22.107 (df = 73)	18.249 (df = 211)	17.836 (df = 211)	17.913 (df = 209)
F Statistic	1.630* (df = 23; 75)	2.065** (df = 23; 75)	1.854** (df = 25; 73)	1.914*** (df = 27; 211)	2.369*** (df = 27; 211)	2.193*** (df = 29; 209)
Note:						*p<0.1; **p<0.05; ***p<0.01

Table 7 shows the results of 6 different specifications for OLS estimation of treatment effects with robust standard errors. For each model the sample includes all dictators who gave more than zero percent in the redistribution phase of the survey. Controls include in this specification are: age, score on CAPTCHA task, education level, income, and gender

Table 8: Red State vs. Blue State - OLS Model

	Dependent variable:					
	percent					
	Red - Race	Red - Effort	Red -Interaction	Blue - Race	Blue - Effort	Blue-Interaction
	(1)	(2)	(3)	(4)	(5)	(6)
white	7.538**		5.395	−3.061		−2.896
	(3.341)		(4.483)	(2.864)		(3.816)
effort		9.174**	6.606		11.535***	12.035***
		(3.523)	(5.077)		(2.629)	(3.846)
whitexeffort			4.506			−0.860
			(6.703)			(5.536)
Observations	172	172	172	211	211	211
R^2	0.226	0.235	0.263	0.163	0.238	0.244
Adjusted R^2	0.094	0.104	0.125	0.045	0.130	0.127
Residual Std. Error	20.985 (df = 146)	20.868 (df = 146)	20.616 (df = 144)	19.072 (df = 184)	18.202 (df = 184)	18.228 (df = 182)
F Statistic	1.706** (df = 25; 146)	1.791** (df = 25; 146)	1.907*** (df = 27; 144)	1.379 (df = 26; 184)	2.206*** (df = 26; 184)	2.095*** (df = 28; 182)
Note:						*p<0.1; **p<0.05; ***p<0.01

Table 8 shows the results of 6 different specifications for OLS estimation of treatment effects with robust standard errors. For each model the sample includes all dictators who gave more than zero percent in the redistribution phase of the survey. Controls include in this specification are: age, score on CAPTCHA task, education level, income, political leanings, and gender

Table 9: Red State vs. Blue State: Selfish Type - Logit Model

	Dependent variable:					
	Likelihood of Being "Selfish" Type					
	Red - Race	Red - Effort	Red -Interaction	Blue - Race	Blue - Effort	Blue-Interaction
	(1)	(2)	(3)	(4)	(5)	(6)
white	−0.111		0.164	0.015		−0.192
	(0.280)		(0.399)	(0.311)		(0.443)
effort		0.304	0.610		0.175	−0.026
		(0.284)	(0.419)		(0.304)	(0.429)
whitexeffort			−0.597			0.410
			(0.574)			(0.620)
Observations	265	265	265	283	283	283
Log Likelihood	−157.616	−157.121	−157.238	−147.265	−147.101	−146.880
Akaike Inf. Crit.	375.233	374.242	376.476	350.530	350.202	353.759
Note:						*p<0.1; **p<0.05; ***p<0.01

Note: Table 9 shows the results of a logit estimation by red state vs blue state of the treatment effects for the subset of dictators that gave zero percent in redistribution. Each coefficient measures the log likelihood of giving zero based off the binary being on or off. Controls include in this specification are: age, score on CAPTCHA task, education level, income, and self-described political leaning

Table 10: Regional Comparison for Dictators who gave more than zero - OLS

	Dependent variable:			
	percent			
	South	Northeast	Midwest	West
	(1)	(2)	(3)	(4)
white	1.980	3.396	−4.388	7.510
	(5.042)	(7.021)	(8.001)	(7.699)
effort	9.975*	6.650	1.109	13.414*
	(5.261)	(8.237)	(8.982)	(6.790)
whitexeffort	−1.573	−2.543	3.893	−10.553
	(7.750)	(12.445)	(10.851)	(9.789)
Observations	153	68	70	92
R^2	0.225	0.170	0.214	0.316
Adjusted R^2	0.057	−0.183	−0.063	0.057
Residual Std. Error	21.153 (df = 125)	20.839 (df = 47)	20.084 (df = 51)	19.690 (df = 66)
F Statistic	1.343 (df = 27; 125)	0.481 (df = 20; 47)	0.773 (df = 18; 51)	1.221 (df = 25; 66)
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 10 shows the results of 4 different specifications for OLS estimation of treatment effects with robust standard errors. For each model the sample includes all dictators who gave more than zero percent in the redistribution phase of the survey. Controls include in this specification are: age, score on CAPTCHA task, education level, income, political leanings, and gender

Figure 1: Distribution of Dictator Payouts

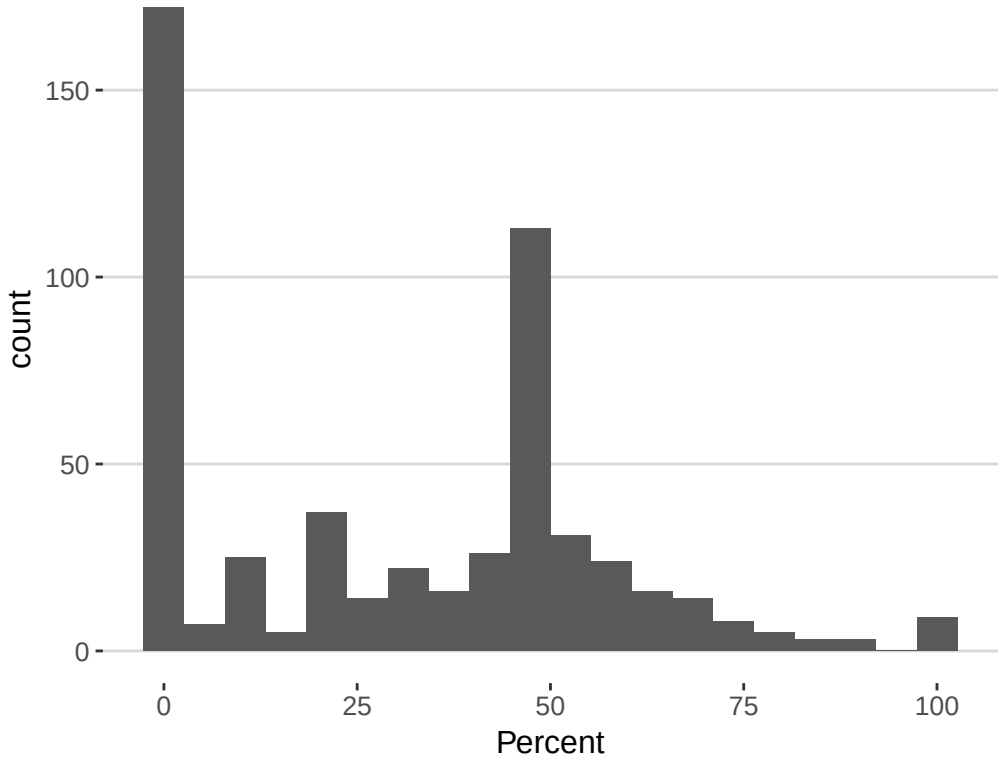

Figure 1a: Distribution of Dictator Payouts by Gender

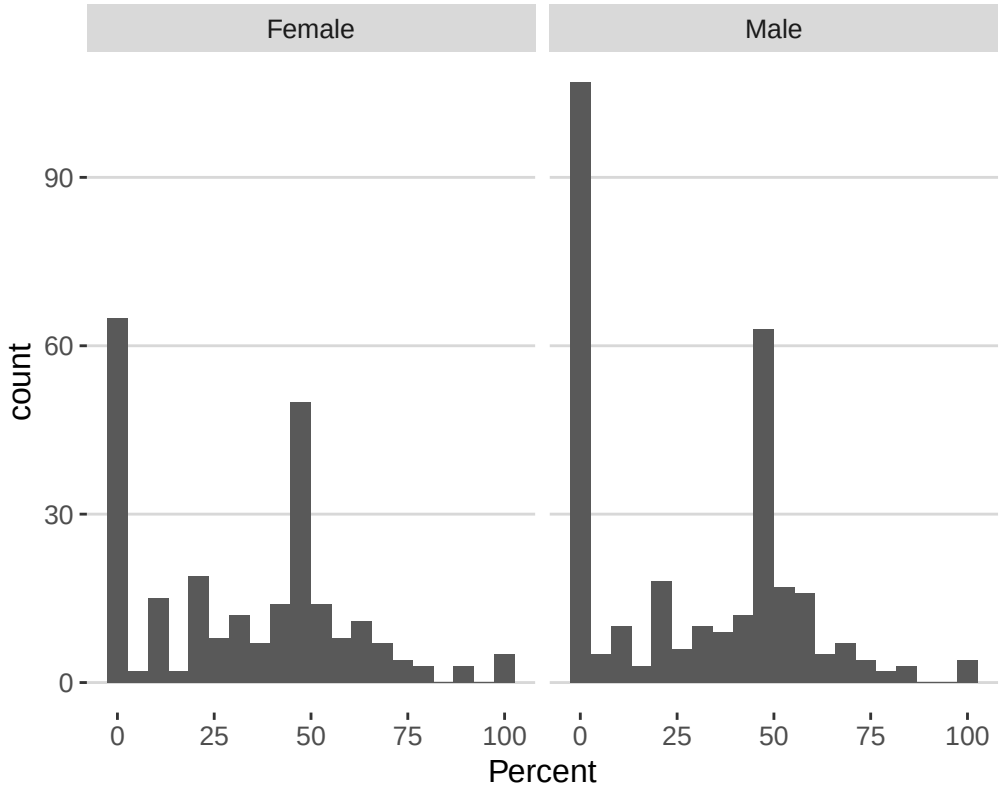

Figure 2: Distribution of Dictator Scores

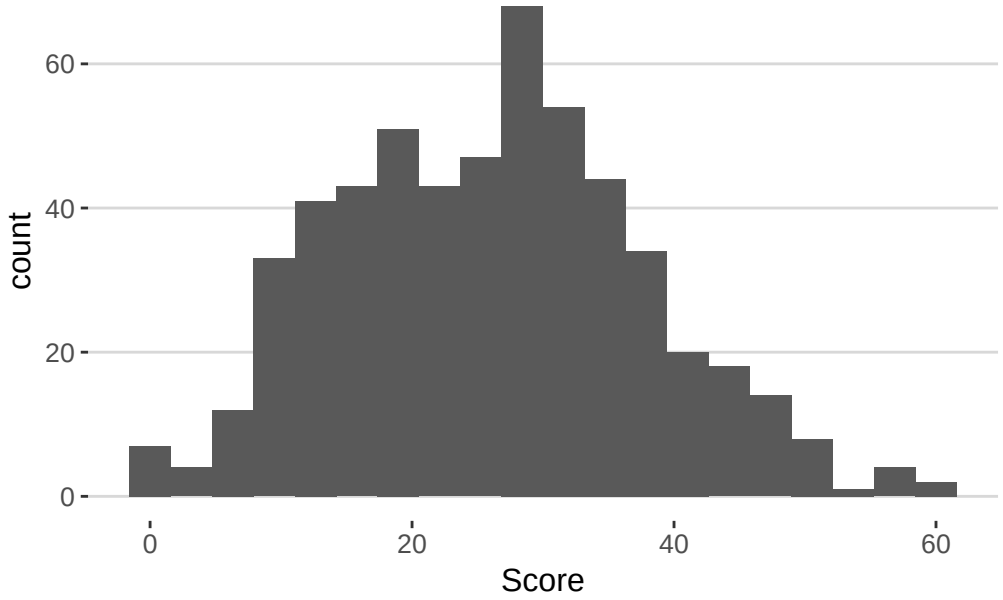

Note: Figure 2 presents a histogram distribution of the scores of dictators on the CAPTCHA data task. Each respondents had 2 minutes to complete as many two word phrases accurately as possible.

Figure 2a: Distribution of Dictator Scores by Gender

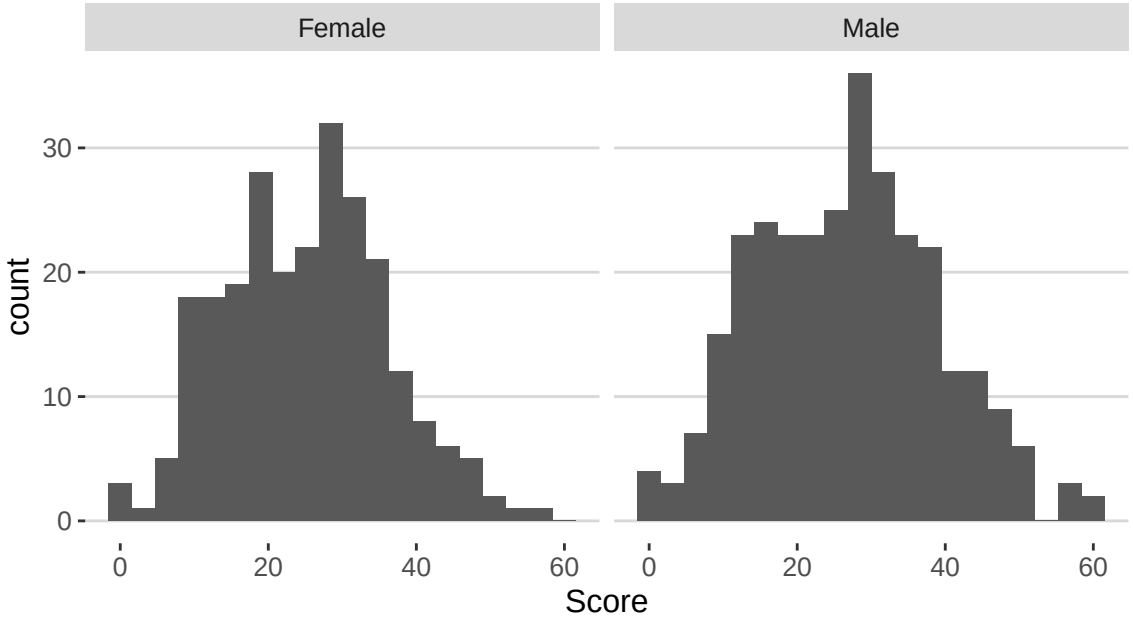

Note: Figure 2a presents a histogram distribution of the scores of dictators on the CAPTCHA data task separated by gender. Each respondents had 2 minutes to complete as many two word phrases accurately as possible.

Figure 3: Distributor Type by Treatment

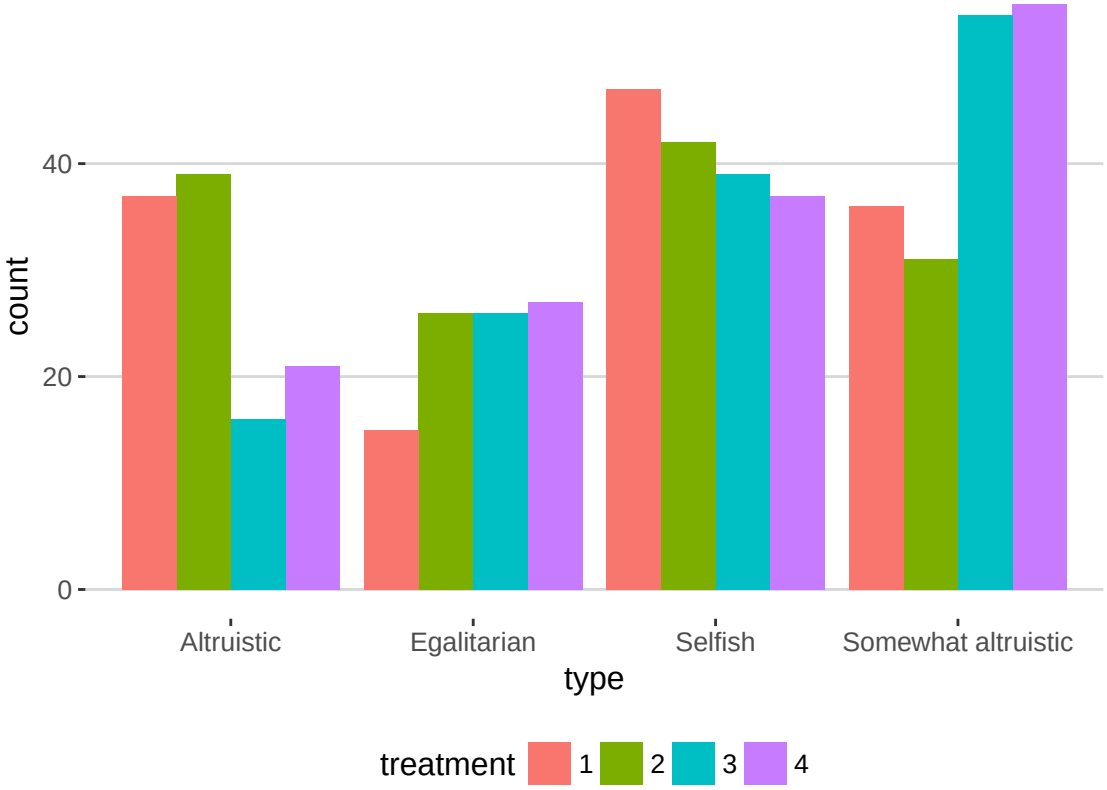

Note: Figure 3 shows a bar graph of the distribution of dictators into different types depending on their pattern of giving in the redistribution task.
Altruists are dictators who gave more than 50% of total income to their partner
Egalitarians gave exact 50% of total income to their partner
Those in the selfish type gave 0% of total income to their parnter.
The somewhat altruists gave greater than 0% but less than 50%

Figure 3a: Distributor Type by Treatment and Gender

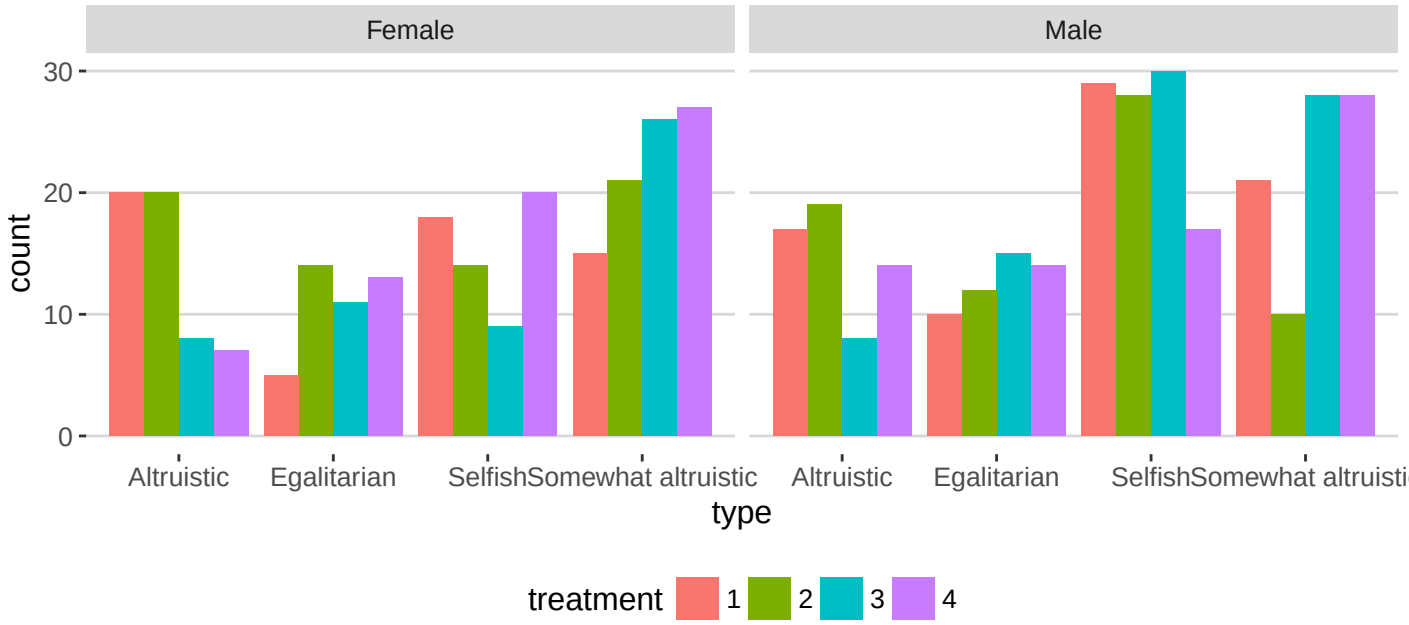

Note: Figure 3a shows a bar graph of the distribution of dictators into different types depending on their pattern of giving in the redistribution task and separated by gender.
Altruists are dictators who gave more than 50% of total income to their partner
Egalitarians gave exact 50% of total income to their partner
Those in the selfish type gave 0% of total income to their parnter.
The somewhat altruists gave greater than 0% but less than 50%

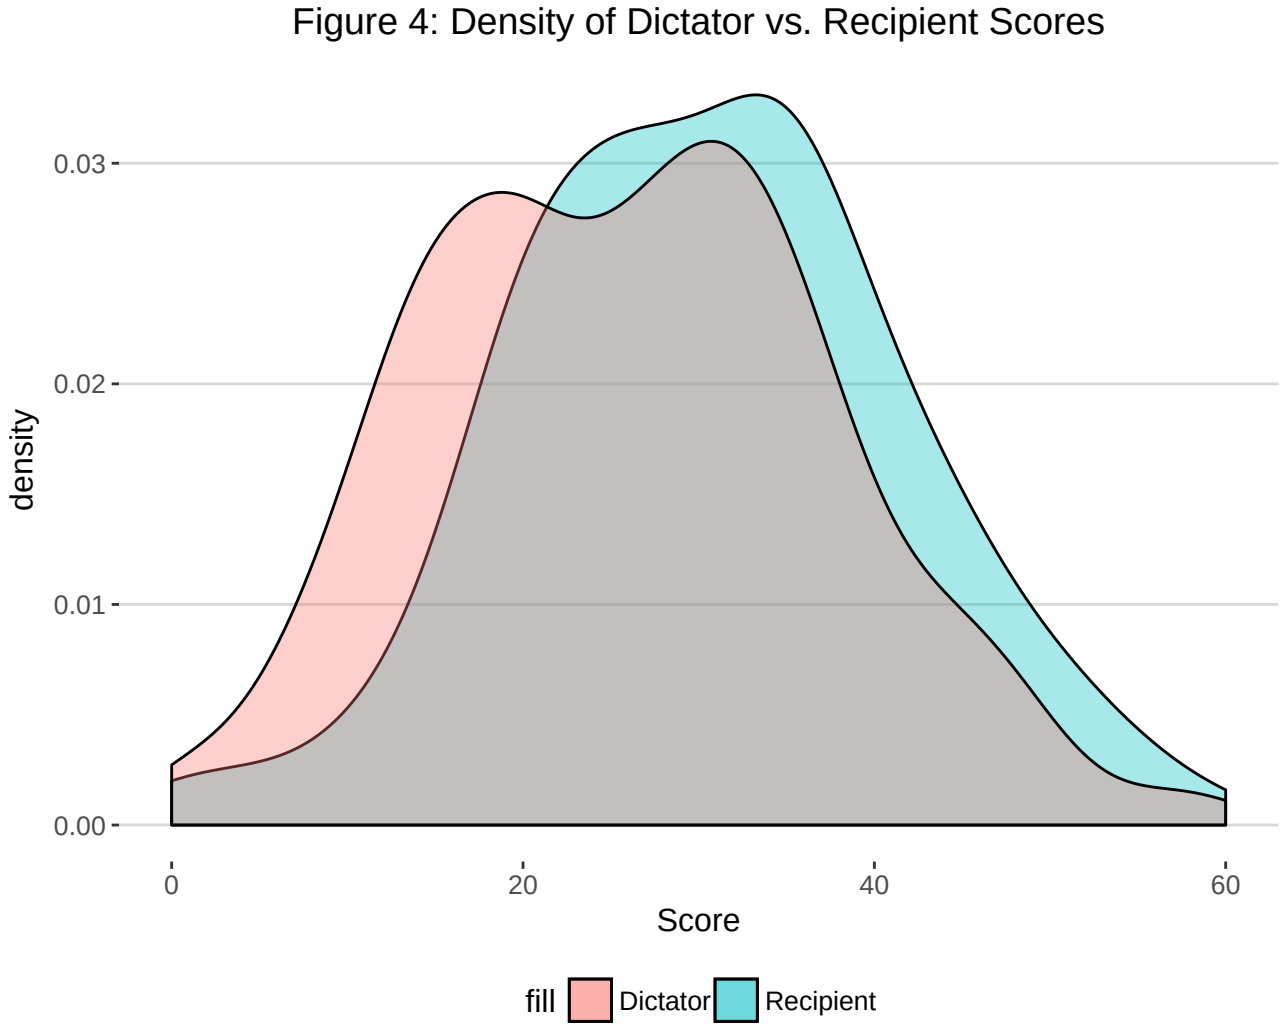

Figure 4 presents a comparison of the density plots of the distribution of scores in the dictator population and the recipient population.

Figure 5: Box Plot Estimates by Treatment

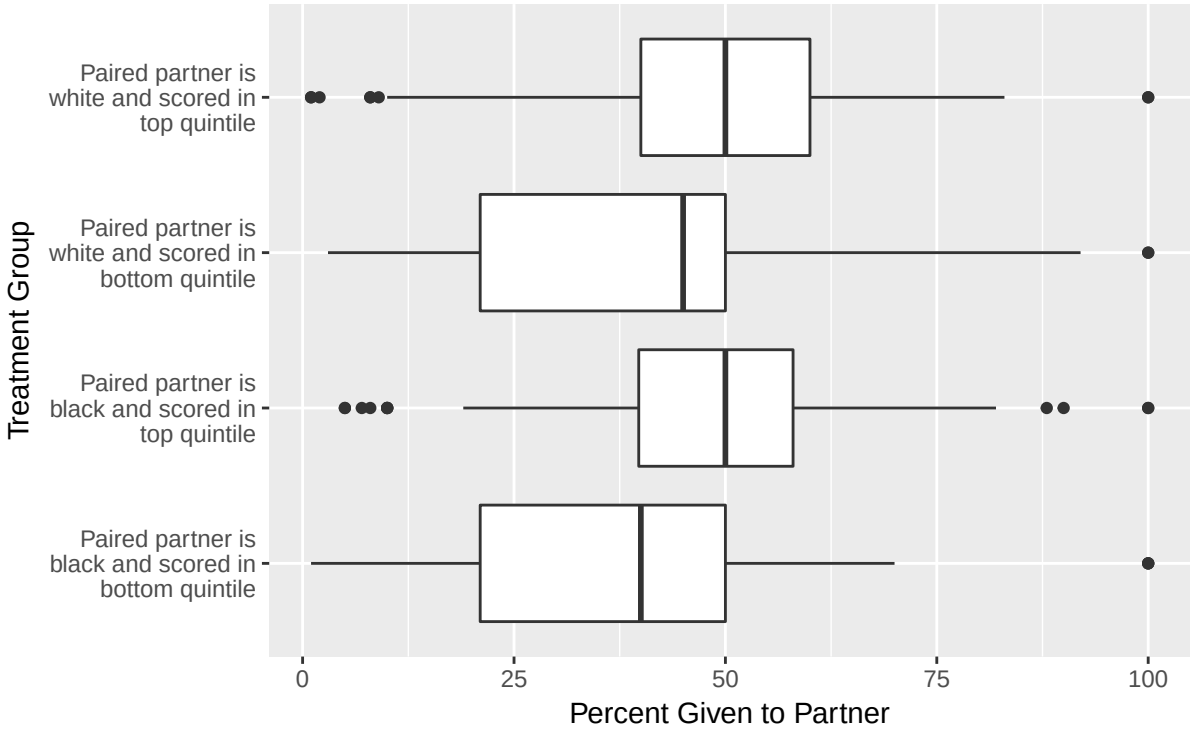

Note: Figure 5 presents box and whisker plots of the four different treatment groups used as interventions treatment assignment dictators were shown a picture of their paired partner's hand and their partner's score on a CAPTCHA data entry task. Respondents who gave zero were excluded from this chart. Each treatment had around 130 participants.

Figure 5a: Box Plot Estimates by Treatment and Gender

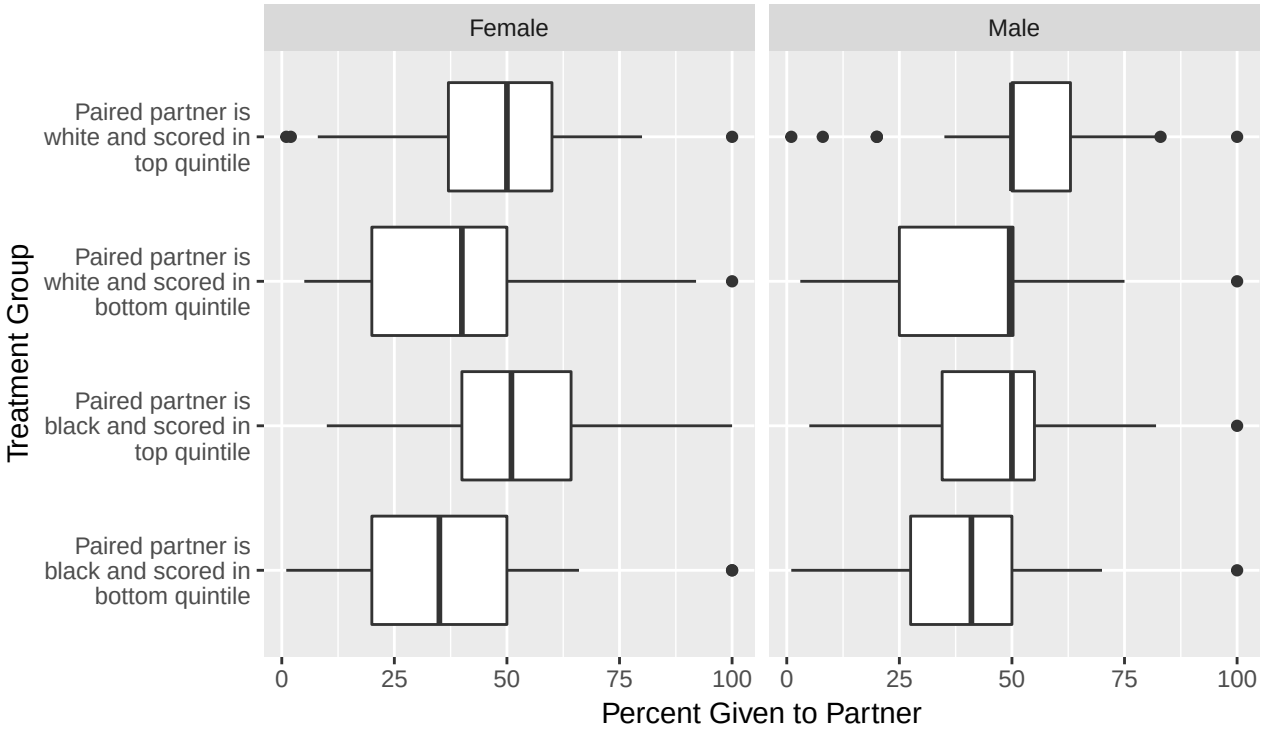

Note: Figure 5a presents box plots by gender of the four different treatment groups used as interventions treatment assignment dictators were shown a picture of their paired partner's hand and their partner's score on a CAPTCHA data entry task. Respondents who gave zero were excluded from this chart. Each treatment had around 130 participants.

8 Appendices

Appendix 1: Sample of CAPTCHAs for data entry task

Appendix 2a: Page 1 of Treatment portion of dictator survey Example 1:

Berkeley
UNIVERSITY OF CALIFORNIA

In the previous task you completed **22** CAPTCHAs.

You have been randomly paired with a partner who also completed the CAPTCHA task. Your partner completed a total of 21 CAPTCHAs.

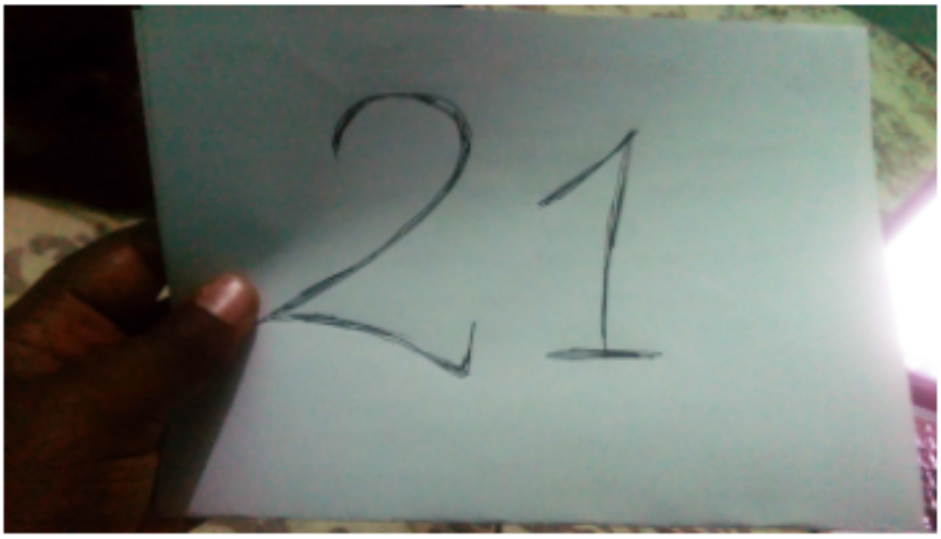

The total value of your combined effort is: **$0.43**

Example 2:

In the previous task you completed **32** CAPTCHAs.

You have been randomly paired with a partner who also completed the CAPTCHA task. Your partner completed a total of 19 CAPTCHAs.

The total value of your combined effort is: **$0.51**

Appendix 2b: Page 2 of Treatment portion of dictator survey

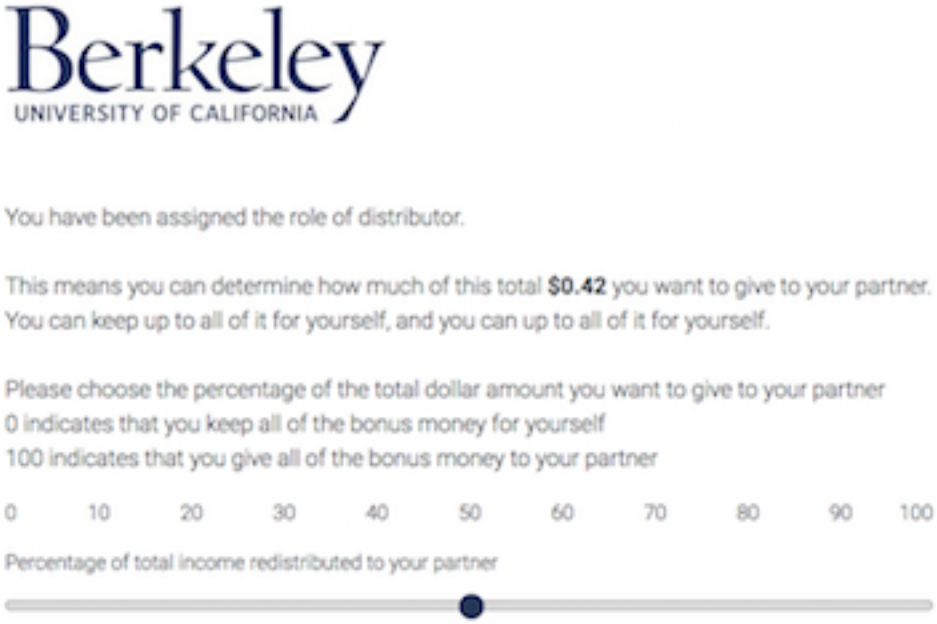

Appendix 3: Language used before CAPTCHA data completion task

On the next page you will have 2 minutes to complete as many Captcha's as you can. To complete a Captcha simply type the two words you see.

You will be paid a bonus of $0.01 for each captcha accurately completed within the minute. Additionally, depending on your performance you may be eligible for future bonuses from a follow-up study.

The timer will start when you continue to the next page.