

Deconstructing Bias: Individual Groupiness and Income Allocation

Rachel Kranton, Matthew Pease, Seth Sanders, and Scott Huettel*
August 2017

Abstract: This paper finds significant, divergent patterns in how people allocate income in group settings. The results indicate that the tendency to favor people conditional on group affiliation, which we call “groupiness,” could be an individual trait. Each participant allocates income in two group treatments, an arbitrary minimal group setting and a political group setting. Many subjects are “not groupy,” showing no favoritism to ingroup in either setting; others are “groupy,” with equally positive favoritism in both. Less than half of subjects are “conditionally groupy,” with greater favoritism in the political group treatment. Using latent class models, we structurally identify nine distinct patterns of behavior. The most prevalent type, 23% of subjects, weighs own and other subjects’ income similarly regardless of group affiliation of others; the second most prevalent type, 20%, puts almost no weight on other subjects’ income regardless of group affiliation of others. Both show no ingroup favoritism albeit in different ways. Twelve percent of subjects’ have particularly high favoritism in both settings. Overall, three of our nine types are not groupy, three are groupy, and three are conditionally groupy. Thus, observed bias in a group setting might not be due to the nature of the setting but rather the selection or composition of individuals within the group.

* Contacts: rachel.kranton@duke.edu, pease.matthew@alumni.duke.edu, seth.sanders@duke.edu, scott.huettel@duke.edu. Authors are listed in alphabetical order following the convention in economics, except for the lab director who is listed last following the convention in psychology & neuroscience. We thank three anonymous referees and the editor for their comments which helped us to hone the analysis and its interpretation. We are grateful for feedback and suggestions from George Akerlof, Jeff Butler, John Miller, Pedro Rey-Biel and seminar and conference participants at Berkeley, Canadian Institute for Advanced Research, Conference on the Economics of Interactions & Culture (EIEF), Duke, Ecole Polytechnique, Erasmus Choice Symposium, Maryland, Paris School of Economics, Pompeu Fabra, Institute for Economic Analysis & Universitat Autònoma de Barcelona, Sciences Po, Stanford, THEEM, and Washington University. We thank Catherine Moon, Robert Richards, and Sierra Smucker for research assistance. The Social Science Research Institute at Duke sponsored our faculty fellows program in 2010-2011, “From Brain to Society (and Back),” and we are grateful for funding from the Transdisciplinary Prevention Research Center (TPRC) at Duke, as supported by the National Institute on Drug Abuse grant DA023026.

I. Introduction

Group conflict is a continual feature of human societies. A long tradition of social psychology studies the implications of social divisions, and this tradition has more recently shaped several areas of experimental economics. Studies largely suggest that people placed in groups favor those within their group.¹ This paper presents a novel within-subject experiment on group divisions and income allocation which deconstructs this bias. Participants allocate income in two separate group settings: minimal groups and political groups. The findings indicate significant heterogeneity in response to the group condition. These results imply that the tendency to favor people conditional on group affiliation, which we call “groupiness,” could be an individual trait.

The results, from the raw data and from the structural estimation, present a new picture of individuals, groups, and identity. For most subjects, allocations are similar in each setting, implying that the group identity per se does not particularly matter; one set of subjects is “not groupy,” never exhibiting bias regardless of group setting, and another set of subjects is “groupy,” always biasing allocations against outgroup members regardless of group setting. Indeed, the subjects that bias their allocations the most do so in both group treatments. The political group does elicit different behavior for some subjects, whom we call “conditionally groupy,” but the response is not uniform; some subjects show mild bias and others show strong bias against political outgroup members.

1. For a broad overview of social psychology on social conflict and groups, see, for example, Berreby (2008). This paper builds particularly on social identity theory (Tajfel & Turner (1979)) and the experimental tradition of Sherif et. al. (1961) and Tajfel et. al. (1971), where in the latter participants are divided into two groups and then perform an experimental task that involves benefits and costs for individuals in the groups. For an historical review of social identity theory see Hornsey (2008); for a meta-analysis of minimal group experiments in the Tajfel et. al. (1971) tradition, see Pechar & Kranton (2017); for further discussion of social psychology experiments and economics, see Akerlof & Kranton (2010) and especially Chen & Li (2009). Section II provides discusses the economics experimental literature.

Hence, observed bias in a group setting might not be due to the nature of the setting but rather the selection or the composition of individuals within the groups.

Overall, our structural analysis identifies nine distinct patterns of behavior. Using the groupiness criteria, we then have a nested framework to understand behavior across the group treatments. Three of the types are not groupy, with the most prevalent type always treating other subjects well, regardless of group treatment and whether recipient is ingroup or outgroup. Three of the estimated types are groupy, always treating the outgroup worse, but to varying extents, from high favoritism in both settings to low favoritism in both settings. The remaining three types are conditionally groupy, with differential utility weights in the political condition and again with varying extents of favoritism. None of the nine types clearly match the average, which appears to be a composite of types. Thus, the analysis indicates not only a wide variety of individual reactions to group settings, but also that the average is not predictive of any individual behavior.

In the experiment, subjects allocate income to themselves and to other subjects in three conditions. The task itself, following Charness & Rabin (1999), is to choose allocations in twenty-six different allocation matrices presented randomly. In the non-group control, subjects allocate income to themselves and to a random participant in the experiment. In the minimal group treatment, subjects are divided into two groups according to answers to a questionnaire on preferences for lines of poetry, paintings, and landscape images. Each subject then (i) allocates income to self and to a recipient in the subject's own group, and (ii) allocate incomes to self and to a recipient in the other group,

presented randomly.² The political group treatment is identical except that subjects are divided into two groups – Democrat and Republican – according to a political questionnaire. The order of the group treatments is randomized.³

The experiment was initially cast to test whether individual subjects' allocations show more bias when they identify more closely with their assigned groups. The minimal group treatment would serve as a control for general group bias, and the political group treatment would differentially affect subjects depending on their party affiliations (Democrat, Republican or Independent). In our previous paper, Kranton, Pease, Sanders, and Huettel (2016), we find that many Democrats show similar bias in the minimal group and the political group treatments, indicating the political treatment has only a small effect beyond a group effect per se for some subjects.⁴ Democratic-leaning Independents, on the other hand, show little bias in both treatments, indicating that group effects are generally small for other subjects.⁵

These findings spur the present analysis which directly estimates individual behavior across the two group treatments. As discussed above, we analyze the raw data and we structurally estimate, via latent class models, the parameters of a parsimonious

2. Separately, subjects allocate income between an ingroup participant and an outgroup participant. We study the data only from the decisions described in the text, where the participant's decisions affected own income.

3. Possible order effects are discussed below in Section V.

4. Using the Fehr & Schmidt (1999) utility function and a latent class model, Kranton, Pease, Sanders and Huettel (2016) estimates four types, categorizes subjects as types, and then cross-tabulates to see if individuals change types depending on group treatment and ingroup or outgroup match. The present paper considers a parsimonious utility function and uses the panel data and full power of latent class models to directly estimate types that characterize behavior in the two group treatments.

5. Both self-declared Democrats and self-declared Democratic-leaning Independents reported similar views on a host of contentious policy issues.

utility function. The objective is to discern any patterns of behavior across the group settings and the prevalence of individuals with these patterns. Selection criteria applied to the latent class models indicate a model with nine types. The nine types, as discussed above, each have distinct patterns of behavior and exhibit different levels of bias or the absence of bias. We then classify individual subjects as types, using subject's actual choices in the experiment to construct posterior probabilities that a subject is a certain type. This classification shows the estimation well captures behavior and allows us to study demographic and other possible correlates of behavior.

The results indicate that standard descriptors do not generally relate to “groupiness.” While many of the tests have low power due to small numbers, the nine types are largely similar in demographics, with two possible exceptions distinguishing groupy types. Fewer women are groupy types than conditionally groupy types, and fewer subjects with highly educated fathers are the most biased groupy type compared to the largest not groupy type.⁷ We find little meaningful differences across the nine types in group-related behaviors outside the lab: religious service attendance, trust in others, and political party affiliation.

Furthermore, we ask whether social preferences in the non-group condition correlate with groupiness in the group treatments. Estimating and categorizing subjects as utility-types in the non-group condition and cross-tabulating, we see only one mapping from these social preferences to groupiness; subjects who put negative weight on others' income in the non-group condition are all groupy in the group treatments. Overall, this

7. About one-third of the subjects have fathers with a beyond-four-year degree, which could proxy for family income.

exercise indicates that groupiness is a different phenomenon than social preferences per se.

The paper is organized as follows. Section II places the paper in the literature, and Section III describes the experiment in detail. Section IV studies favoritism in the raw data and defines our notions of groupiness. Section V presents the latent class models, the results, and interprets the structural estimation of the utility-types. Section VI studies demographic and social preference correlates of the nine types. The Conclusion summarizes the findings and suggests directions for further research.

II. Advancing the Literature

This experiment and analysis advance three strands of research. First, this paper advances our understanding of identity and economic choices. Identity here, as in social psychology, describes an individual in terms of a social category or group, such as gender, race, ethnicity, nationality, political party (Akerlof & Kranton 2000, 2010). One experimental approach to studying the impact of identity employs such “natural groups,” with findings, for example, that the race, ethnicity, or political party of subjects relates to play in dictator and ultimatum games (e.g., Fershtman & Gneezy (2000), Glaeser, Laibson, Scheinkman, & Souter (2000), Fowler & Kahn (2007)).⁸ Following the social psychology literatures, a second method creates social categories inside the laboratory, as in Chen & Li's (2009) minimal group experiment on social preferences.⁹ The present

8. Further work shows that natural groups impact play in prisoner's dilemma, public goods and trust games (e.g., Goette, Huffman & Meier (2006), Bernard, Fehr, & Fischbacher (2006)). In an experiment studying redistribution, Klor & Shayo (2010) divide subjects according to their university fields of study and find subjects vote more often for the tax rate that favors ingroup members.

9. Other economic experiments using arbitrary groups, with different tasks, include Charness, Rigotti & Rustichini (2006), Chen & Chen (2011) and Hargreaves Heap & Zizzo (2009).

study uses both methods in a within subject design; the design then allows us to see the same individuals in different group settings and ask whether the nature of the group division matters or whether groups per se are important (or not).¹⁰

Second, the paper contributes to the experimental study of income allocation by uncovering patterns in bias in group settings. Much of the literature on which we build estimates “social preferences,” which represent the relative value people place on their own and others’ incomes, often depending on whether others’ incomes are higher or lower than one’s own income.¹¹ The literature largely concludes that on average subjects are inequity averse or seek to maximize total income (e.g, Charness & Rabin (2002), Camerer & Fehr (2004)). Another track of this literature emphasizes individual heterogeneity in social preferences (e.g., Bolton & Ockenfels (2000), Andreoni & Miller (2002), Engelmann & Strober (2008), Fisman, Kariv and Markovits (2007)). Many subjects are inequity averse or maximize total income, but many subjects simply seek to maximize own income and some destroy the income of others (e.g. Levine (1998), Fershtman, Gneezy and List (2012), Iriberry and Rey-Biel (2013)). Experiments on social preferences in groups follow both tracks. In seminal work using a minimal group paradigm, Chen & Li (2009) finds that subjects on average are inequity averse and suffer less from advantageous inequality and more from disadvantageous inequality vis-à-vis outgroup participants.¹² Using real world groups, Klor & Shayo (2010) and Kranton,

10. Goette, Huffman & Meier (2012) compare two sets of subjects, one randomly assigned to minimal groups and the other randomly assigned to groups that involve real social interactions leading to social ties. The present paper employs a within subject design and estimates individual patterns across contexts.

11. For a critical review of the social preferences literature see Levitt & List (2007).

12. Kranton, Pease, Sanders, and Huettel (2016) replicates these results in the data from the present experiment.

Pease, Sanders, and Huettel (2016) find heterogeneity of social preferences toward others within a group treatment. The present paper uncovers a new and different sort of individual heterogeneity: responsiveness to group treatments in general.

Finally, this paper deploys methods that are relatively new to experimental economics, to study panel data generated by a within subject design. Several papers in experimental economics have used latent class models to discern types in single experimental conditions.¹³ Latent class models were formulated by criminologists to sift through a myriad of individual choices across time to uncover archetypical criminal behavioral patterns.¹⁴ We exploit the power of this method using our panel data to characterize any archetypical behavioral patterns across experimental conditions.

III. Description of Experiment¹⁵

The experiment was conducted at Duke's Human Neuroeconomics Laboratory, which follows the experimental economics protocol of no deception. The experiment involved 141 subjects drawn from the Duke University community.¹⁶

13. For example, Stahl & Wilson (1995), Stahl (1996), Bosch-Domènech et. al. (2010) estimate the proportion of subjects who reason at different levels. Harrison & Rutström (2009) and Conte, Hay, and Moffatt (2011) allow a mixture of expected utility and prospect theory. Fischbacher, Hertwig, and Bruhin (2013) use a mixing model and classify subjects into types by posterior distributions as in the present paper. Their goal is to study the relationship between response time and play in dictator games as a window on individual heterogeneity of social preferences. Bruhin, Fehr, and Shunk (2016) estimate a latent class models in two identical experiments separated in time and compare the results to find temporal "stability" in social preferences.

14. Nagin (2005), for example, uses arrest data to determine which individuals are likely to become career criminals and which only commit crimes as adolescents. As another example with "real-world" panel data, Bruhin et.al. (2015) use a latent class model to discern canonical behavioral responses over time to a blood-donation policy intervention.

15. The description of the experiment is largely identical to the description in Kranton, Pease, Sanders, and Huettel (2016).

16. Seventy-six percent of the subjects were Duke students, 11% students from other schools (mostly University of North Carolina, Chapel Hill), and the remainder were non-students (mostly staff). Of the

Instructions	3-5 Minutes
<u>Non-Group Condition</u>	
52 Choices	12 Minutes
<u>Minimal Group or Political Group Treatment</u> (randomized)	
Survey	2-5 Minutes
78 Choices	17 Minutes
<u>Minimal Group or Political Group Treatment</u> (randomized)	
Survey	2-5 Minutes
78 Choices	17 Minutes
Post Experiment Survey	3-5 Minutes

Figure 1. Timeline of Experiment

Sessions proceeded as in Figure 1. Subjects received instructions on the decisions they would make and practiced using the computer keys that would indicate their choices. (See the Appendix for instructions.) All sessions began with the *non-group* condition. Each subject then made decisions in the *minimal group treatment* and the *political group treatment*, with the order randomized across subjects. The post-experiment survey asked for demographic information (e.g., age, sex, major field of study, hometown) and personal attitudes and practices (e.g., frequency of religious service attendance).

In the non-group condition, subjects allocated money to themselves and randomly selected participants in two kinds of pairings: (1) themselves and other subjects, labeled

students, 86% percent were undergraduates. Eighteen percent of all subjects were born abroad. Sixteen percent were born in North Carolina, 12% in New York or New Jersey, and 6% in California, with the rest of the subjects born in one of 28 states or the District of Columbia. Students reported a wide range of major fields of study, many listing multiple fields. In all, 27 different fields were mentioned, with the most mentioned as follows: biology 21%, psychology/neuroscience 16%, economics 8%. The pool was 65% female.

YOU-OTHER, and (2) between two other subjects, labeled OTHER-OTHER.¹⁷ The screens indicated the pairing, as in Figure 2 below for YOU-OTHER. The pairings occurred randomly. For shorthand below, the initials NG (non-group) designate this condition.

In each group treatment, subjects were divided into two groups according to answers to survey questions. In the *minimal group treatment*, subjects were presented pairs of lines of poetry, landscape images, and abstract paintings (by Klee or Kandinsky) and asked which item in each pair they preferred. The items were matched (e.g., the landscape images were almost identical) so that this choice is unrelated to individual subject characteristics. The online Appendix provides examples. Subjects were then divided based on their answers to these questions and were given (true) information about similarity, or not, in answers to survey questions.¹⁸ Subjects then allocated money in three kinds of pairings, presented randomly: (1) between themselves and one own-group member, labeled YOU-OWN, (2) between themselves and one other-group member, labeled YOU-OTHER, and (3) between one own-group member and one other-group member, labeled OWN-OTHER. For shorthand below, we refer to pairings (1) and (2) as MG Ingroup and MG Outgroup, respectively.

17. The latter allocations do not affect a subject's own payoffs. The present paper does not use data from these Other-Other pairings or the Own-Other pairings in the group treatments mentioned below.

18. The online Appendix describes the procedure and the information subjects received about the other participant's answers to survey questions in the minimal group and the political group treatment. In all other ways the matching is anonymous, and the recipient could be from another session of the experiment.

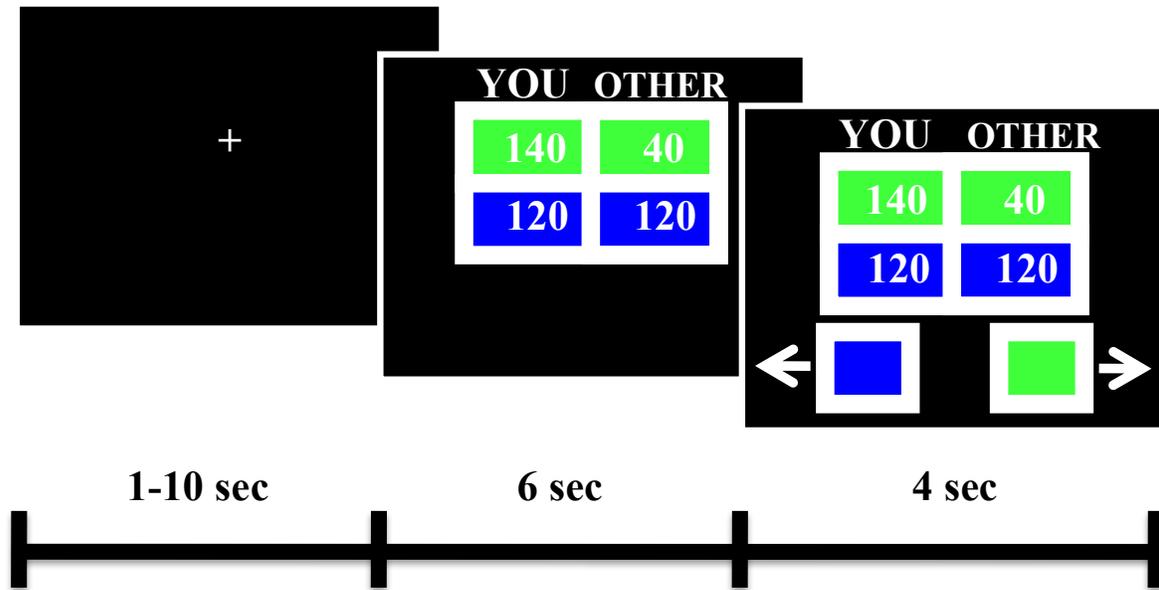


Figure 2. Timing and Presentation of Allocation Choices

The *political treatment* began with a political survey. Subjects were first asked their affiliation as Democrat, Republican, Independent, or None of the Above. The next question asked subjects to refine their political leanings: “strong” or “moderate” for party affiliates, “closer to Democratic” or “closer to Republican” for Independents and None of the Above.¹⁹ Subjects were then asked their opinions on five issues dividing the political spectrum in the United States at that time, as well as on media outlets and religious service attendance. Subjects were then placed into the Democrat group (containing all

19. The Appendix provides the breakdown of the subjects by political party and leanings according to the political treatment survey. Just under half are Democrats (48%) and only 13% are Republicans. Independents and None of the Above make up more than one third of subjects (39%). Of these subjects, 62% are Democratic-leaning.” The subject pool appears to be representative of the Duke University community. Overall the majority (by at least 10 percentage points) of North Carolina’s population is Democratic or “leans” Democratic, with a concentration of Democrats in the region where Duke is located (<http://www.gallup.com/poll/114016/state-states-political-party-affiliation.aspx>). Nationally this age cohort is largely Democratic (<http://www.people-press.org/2011/11/03/the-generation-gap-and-the-2012-election-3/>). The distribution of our subject pool also matches the political spectrum of undergraduates at Princeton, which has a similar undergraduate program and is the one peer institution for which we could find survey data (<http://www.dailyprincetonian.com/2008/11/04/21969/>).

Democrats and “closer to Democratic” subjects) or the Republican group (containing all Republicans and “closer to Republican” subjects). Subjects were given (true) information on similarity and differences in answers to survey questions for their assigned group and for the subjects to whom they allocated income. Subjects allocated income in three types of pairings, YOU-OWN, YOU-OTHER, and OWN-OTHER, with exactly the same format as in the minimal group treatment. Below for shorthand, we refer to the relevant pairings as POL Ingroup and POL Outgroup.

For each kind of pairing in each condition, subjects were randomly presented 26 different 2x2 allocation matrices.²⁰ The Appendix provides the collection of matrices, and Figure 2 provides an example. The rows within each matrix were randomized, and the colors of the rows (blue or green), as well as the left and right keys, were all randomized.

Consider i 's choice in a normalized matrix $\begin{bmatrix} \pi_i & \pi_j \\ \pi_i & \pi_j \end{bmatrix}$, where i earns weakly more in the

top row than the bottom. A subject i that consistently chooses the top row is arguably concerned only about his own income. A subject i that chooses the bottom row is also concerned about j 's income; i sacrifices own income in order to increase or decrease j 's income. The design allows us to study the difference in the decisions i makes when recipient j is in i 's own group vs. when recipient j is in the other group and whether the specific group treatment matters.

In addition to the show-up fee of \$6, subjects received payment for one choice selected at random from each of the three conditions—non-group, minimal group, and

20. These matrices were constructed following Fehr & Schmidt (1999) and Charness & Rabin (2002). The matrices capture subjects' tradeoffs between own income and another person's income.

political group. Following the protocol of the lab, the choices were translated into dollars, and subjects earned about \$15 for a one-hour session.

Before analyzing the data, we discuss possible experimenter demand effects. Subjects might think experimenters are emphasizing groups and act according to what they think experimenters expect. There are several responses to this concern. First, real-world actors create, highlight, and exploit group divisions, and the aim of this experiment, following a long tradition in social psychology, is to see how people behave in such circumstances. Second, if there is a demand effect, there is apparently no common understanding as to what the demand is; we find most subjects do not differentiate between ingroup and outgroup, and among those that do, there is heterogeneity in behavior. Finally, some might argue that the political group treatment would have a higher demand effect, but the political treatment is also more salient by design. If there is such a differential, again there is little commonality among subjects as to what the differential implies.

IV. Income Allocations in the Raw Data

This section studies the differences in income allocations to ingroup and outgroup participants in the raw data.²¹ Consider the following individual measure of bias: For each subject i , in each group treatment g , take each matrix m faced by agent i , $m = \{1, \dots, 26\}$, and the choice of π_j when j is in i 's own group versus when j is in the other group. The difference, $\Delta_i(m)$, is positive when i gives more to the subject in his

21. In addition to the study of the raw data below, we conducted a factor analysis of subjects' choice data which shows (1) subjects make consistent choices on matrices that are shown, by the analysis, to be similar, (2) subjects have heterogeneous choice patterns, and (3) subjects are sensitive to the losses in own income when choosing allocations. The model and analysis are available upon request.

group for matrix m . In each group treatment g , for each subject i the average of these differences yields an individual statistic we call (ingroup) *favoritism*: $d_i(g) =$

$$\frac{1}{26} \sum_m \Delta_i(m).^{22}$$

The maximum possible favoritism is 69.23.²³

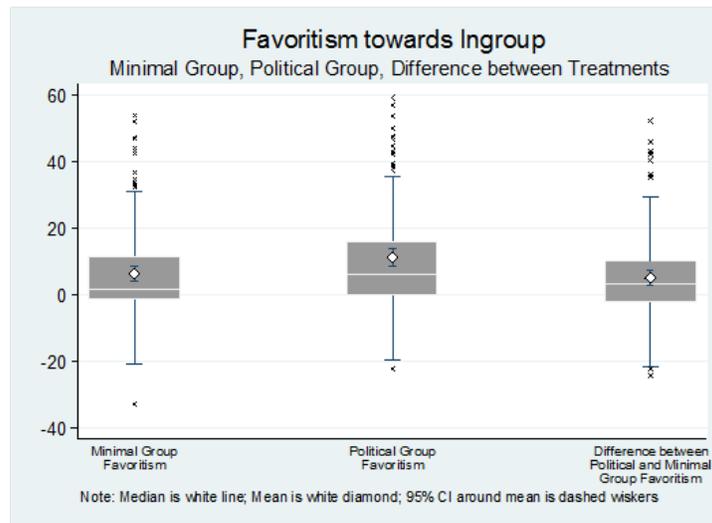


Figure 3.
Distributions of Favoritism: Minimal Group, Political Group, and Difference

The distributions of favoritism, shown in Figure 3, exhibit a wide range of bias in the group treatments. The two left-hand-side box-and-whisker plots in Figure 3 provide the medians, interquartile ranges, and outliers for favoritism in the minimal group and for favoritism in the political group. Superimposed are white diamonds for the means and the 95% confidence intervals around the means. In the minimal group, both the median

22. Another measure of favoritism in the raw data could be how much individuals give up of their own income in order to increase or decrease the payoffs of ingroup versus outgroup participants, or the ratio of income. In previous analyses of the data (Kranton, Pease, Sanders, and Huettel (2013), we find that participants are price-sensitive in that they more often choose allocations that involve little decrease in own payoffs. In the online Appendix for present paper, we provide the differences in i 's own payoffs in ingroup versus outgroup matches; these differences are small but do vary across the subject population and mostly track the favoritism measure defined in the text.

23. This amount is calculated by subtracting, for each matrix, the lowest possible income for j from the highest possible income for j .

and mean are positive, 1.6 and 6.28 respectively, with the mean significantly different than zero (at the 0.01 level) but the median is not statistically different than zero. The mean is higher than the median, pulled up by the many outliers who allocate between about 30 and 55 more to ingroup participants than to outgroup participants. The political condition shows a similar pattern, shifted up slightly; the mean, at 11.31 is higher than the median of 6.15, and outliers range from 35 to about 60.²⁴ The right-hand-most plot in Figure 3 shows the distribution of the difference in minimal group and political group favoritism levels and shows that at least half of subjects have a difference close to zero, with a median of 3.2 and mean of 5.03.²⁵ A set of outliers favors the ingroup much more in the political group than in the minimal group.

The question then becomes whether there is a correlation between favoritism in the minimal group and favoritism in the political group, e.g., whether the participants who are outliers in the minimal group treatment are the same as the participants who are outliers in the political group and hence are near the median in the right-hand-most plot.

Figure 4 below illustrates our main finding; individual favoritism in the political treatment is strongly correlated with favoritism in the minimal group setting. Figure 4 plots each subject's favoritism in the minimal group on the x-axis and favoritism in the political group on the y-axis. The dashed line is the 45° line. The data appear to show a strong correlation and indeed the correlation coefficient is 0.63; the linear regression (not shown) has an R^2 of 0.4.

24. This median and mean are each significantly greater than zero at the 0.01 level.

25. This median and mean are each significantly greater than zero at the 0.01 level.

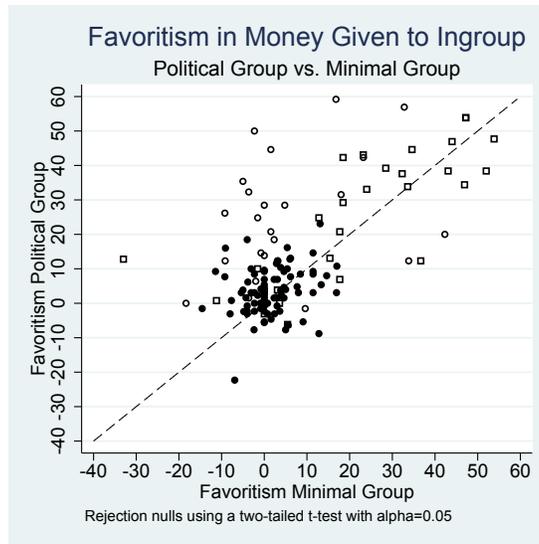


Figure 4.
Individual Favoritism Levels: Political Group vs. Minimal Group

To interpret the correlation in the data, we introduce the notion of *groupiness*. The following definitions are parsimonious and can be applied to raw data in any experimental task where individuals allocate income to participants in different groups.²⁶ Using statistical tests on each individual’s favoritism measures, we divide the subjects into three sets, which we define as follows, rejecting the null hypotheses at the 5% level using two-tailed t-tests. A subject is *not groupy* if we fail to reject zero favoritism in the minimal group treatment and we fail to reject zero favoritism in the political group treatment. A subject is *groupy* if we fail to reject that favoritism is zero in either the minimal group or political group and for whom we fail to reject that the difference in the favoritism levels is equal to zero. A subject is *conditionally groupy* if we fail to reject

26. We developed this definition to be both meaningful and easy to implement and interpret. We entertained several alternative definitions of “groupiness,” including measures based on the ratio of the own to others’ payoffs and measures based on the utility function parameters. The alternatives yield similar qualitative results but are harder to estimate and harder to interpret.

that favoritism is zero in either the minimal group or political group and for whom we reject equal favoritism. In Figure 4, not groupy subjects' data are shown as black dots, and, by the definition, are clustered around (0,0). Sixty-eight percent of subjects are not groupy. Groupy subjects' favoritism levels are indicated by squares in the figure, and, by the definition, track the 45° line. Groupy subjects consist of 19% of the subjects. Conditionally groupy subjects' favoritism levels, shown as open circles, are mostly above the 45° line. They are the smallest subset of our subject pool consisting of 13% of the subjects.²⁷ Given the null hypotheses, the low power of tests on each individual's favoritism levels (which are based on only 52 decisions each) could inflate the number of not groupy subjects. We address this problem below with the latent class model, which pools the data and identifies types of subjects.

The dispersion of subjects along and above the 45° line also suggests that subjects are making systematically different allocation decisions. While the favoritism measures and groupiness designations summarize bias against the outgroup, these outcomes could emerge in different ways. A not groupy subject, for example, could be giving nothing to all other subjects regardless of group, or giving as much as possible to all other subjects regardless of group. A groupy subject could be giving much more to the outgroup than the ingroup in both treatments or just a bit more to the outgroup than the ingroup in both treatments. The structural analysis of the next section estimates utility functions using latent class models to characterize possibly distinct patterns of behavior across the two group treatments and test whether utility parameters are the same or not vis-à-vis ingroup

27. The online Appendix provides a cross-tabulation between individuals' groupiness designations based on individual data alone and individuals' groupiness designations based on categorization as a type in the latent class model. We discuss this cross-tabulation below.

and outgroup participants in the two treatments. With these estimated types, we can also return to the study of groupiness, based on estimated favoritism levels of each type.

V. Structural Estimation of Social Preferences

This section posits and estimates a parsimonious utility function to discern patterns in behavior across group settings, if any. Suppose an individual i 's utility is a simple additive function of own and the other's income:

$$U_i(\pi_i, \pi_j) = \beta_i \pi_i + \rho_i \pi_j,$$

where β_i is the value i places on own income and ρ_i is the value i places on person j 's income. In a group setting, the values i places on π_i and π_j could depend on the nature of the group division and whether j is in i 's group or not. We consider utility

$$U_i(\pi_i, \pi_j; g, J) = \beta_i(g, J) \cdot \pi_i + \rho_i(g, J) \cdot \pi_j$$

where $g = \{MG, POL\}$ again indicates the particular group division and $J = \{In, Out\}$ indicates whether j is in i 's group or not, yielding eight utility function parameters:

$\beta_i(MG, In)$	$\rho_i(MG, In)$
$\beta_i(MG, Out)$	$\rho_i(MG, Out)$
$\beta_i(POL, In)$	$\rho_i(POL, In)$
$\beta_i(POL, Out)$	$\rho_i(POL, Out)$

In most experiments, individual-specific parameters cannot typically be estimated, since each subject would need to make more decisions than is feasible in an experimental setting to yield precise estimates.²⁸ In our setting, individual estimates are complicated

28. To overcome these difficulties and to see patterns in the behavior, researchers studying social preferences have calibrated the extent to which individual utility functions match canonical forms (Andreoni & Miller (2002), Fisman, Kariv, & Markovits (2007)).

by the binary choices facing subjects, and with eight parameters many tests are still low-powered despite the large number of decisions.²⁹

To overcome these difficulties and, more substantively, to discern individual patterns across the group settings, we exploit the panel data of our design and estimate a series of latent class models. Each model posits a number of *types*, T , where each particular type t has a unique set of the eight function parameters (β_t, ρ_t) , and each type t is a proportion of the population p_t , where $\sum_t p_t = 1$. The parameters (β_t, ρ_t) and the proportions are estimated to maximize a likelihood function. To select among these models, we use the BIC and the AIC criteria, which balance the increase in the likelihood function from estimating more types against a penalty for the added model complexity of more parameters.³⁰

Formally, we build our analysis as follows. If each individual's type were known, we could estimate a binary choice model for choosing the bottom row in each matrix for individuals of type t . Assuming an extreme value distribution for the error terms, the parameters could be estimated for the type t individuals by maximizing

$$L(\beta_t, \rho_t) = \prod_{i=1}^{\tau} \prod_{m=1}^{26} \Lambda_{mi}(\beta_t, \rho_t | \pi_i, \pi_j)^{d_{mi}} \left(1 - \Lambda_{mi}(\beta_t, \rho_t | \pi_i, \pi_j)\right)^{1-d_{mi}}$$

where

29. In the matrix task, many participants rarely choose a row with a lower income for self, so the parameter values that characterize the tradeoff of money for self vs. money for others are poorly identified for these individuals.

30. There is little consensus on the best selection criterion (see, e.g., Burnham & Anderson (2004)). We use the BIC and AIC in combination, as discussed below, to determine the model to present in this paper. Other selection techniques, such as estimating a model on a subset of the sample and then testing that model's out-of-sample predictions on another subset of the sample (e.g., Bruhin et. al. (2015)), is not possible with our limited-numbers subject pool.

$$\Lambda_{mi}(\beta_t, \rho_t) = \exp(V_{mi}^{bot} - V_{mi}^{top}) / (1 + \exp(V_{mi}^{bot} - V_{mi}^{top}))$$

and

$$(V_{mi}^{bot} - V_{mi}^{top} | \beta, \rho_t) = \beta_t(\pi_{i,m}^{bot} - \pi_{i,m}^{top}) + \rho_t(\pi_{j,m}^{bot} - \pi_{j,m}^{top}).$$

Since we do not know each individual's type, we condition on an individual being a type and then sum over the distribution of types. That is, for T types, we estimate

$$L(\beta, \rho, p) = \prod_{i=1}^{141} \prod_{m=1}^{26} \prod_{t=1}^T p_t \Lambda_{mi}(\beta_t, \rho_t | \pi_i, \pi_j)^{d_{mi}} (1 - \Lambda_{mi}(\beta_t, \rho_t | \pi_i, \pi_j))^{1-d_{mi}}$$

where $p=(p_1, \dots, p_T)$ is estimated along with the vectors of utility parameters, (β, ρ) , for the T types.³¹ Since type numbering is arbitrary we denote type 1 as the type with the highest fraction in the sample; the T^{th} type has the lowest fraction in the sample.

We estimate each model of $T = \{1, \dots, \tau\}$ types and each time check the BIC and AIC criteria. The BIC criterion, values of which are reported in Table A3 in the Appendix, is essentially tied at eight and nine types, and the AIC criterion suggests more than nine types. Since the nine-type model gives more detail by adding a new type (rather dividing one type into two types) and pulls the selection in the direction of the AIC, we report the nine-type model.³²

31. To insure that $0 \leq p_t \leq 1$ for all t , the mixing distribution is specified as a logistic function with a constant. That is, $T-1$ constants $\{\theta_1, \theta_2, \dots, \theta_{T-1}\}$ are estimated, and the proportion of type 1 is then calculated as $\exp(\theta_1)/(1+(\exp(\theta_1)+\exp(\theta_2)+\dots+\exp(\theta_{T-1})))$ and similarly for the proportion of each type t .

32. In addition, we checked whether the 8, 9, or 10 types models led to different categorizations of subjects to types (categorization described below). In general, parameter estimates for types were extremely close when subjects were not reclassified as the number of types expanded. One exception is that types 7 and 9 in the nine-type model were combined into a single type in the eight-type model. These two types appear distinct in their behavior, an additional reason we settled on the nine-type model. When we estimated the ten-type model, the additional type was poorly estimated and not behaviorally distinguishable from existing types.

Classifying individual subjects as types, using subject's actual choices in the experiment, shows the estimation indeed well captures behavior. Having estimated the model, it is straightforward to calculate the *posterior* probability that a particular subject i is type t . Under the estimated parameters and given the choices that i actually made, the probability of making those choices if i is type t is

$$\Gamma_t(\beta_t, \rho_t) = \prod_{k=1}^{26} \Lambda_{tk}(\beta_t, \rho_t | \pi_i, \pi_j)^{d_{ki}} \times \left(1 - \Lambda_{tk}(\beta_t, \rho_t | \pi_i, \pi_j)\right)^{(1-d_{kt})}$$

Using Bayes' rule with the estimated mixing proportions p_t as priors of being type t , the posterior probability that i is type t , P_t is just

$$P_t(\beta, \rho) = \frac{p_t \Gamma_t(\beta_t, \rho_t)}{\sum_{t=1}^9 p_t \Gamma_t(\beta_t, \rho_t)}.$$

We then categorize individuals as type t based on their posterior probability of being type t . In particular, we assign i type t if $P_t = \max(P_1, \dots, P_9)$. Of the 141 subjects in our experiment, 128 are assigned to their type with probability at or above 0.99. Only ten subjects have a posterior probability below 0.90. These ten subjects are dispersed among types 2, 3, 4, and 6; hence, no single type absorbs these participants. Finally, all of the subjects categorized as types 7, 8, or 9, have posterior probabilities of 0.99 or 1, indicating that these types, while small fractions of the population, well portray these subjects' distinct patterns of behavior.

Table 1 presents the results of the nine-type model. The parameters for the "average subject," which are estimated from a degenerate model with one type, are presented in the last column. The top half of the table reports the parameter estimates for the minimal group; the bottom half contains the political group parameters by providing the difference between the parameters in the minimal group and the political group for

each type, along with tests of whether the differences in the parameters are statistically significant.

An overview of the results shows that Type 1 and Type 4—an estimated 40% of total subjects—each have utility parameters that are statistically the same in the political group and minimal group treatments.³³ The two types, however, are quite different in their weights on outgroup incomes. Type 1 puts similar weights on own income, ingroup income, and outgroup income; Type 4 puts negative weight on outgroup income. The rest of the types have utility function parameters that are statistically different in minimal group and political group, but the implications of the parameters for allocations and the economic significance of these differences is not readily apparent.

To interpret the results in depth, we evaluate both the parameter values and the predicted amounts each type would allocate to ingroup and outgroup recipients in each treatment. The latter are reported in Table 2. From these amounts, we also calculate predicted favoritism levels for each type, also reported in Table 2.³⁴ Using the groupiness definitions in the previous section, we evaluate each type as not groupy, conditionally groupy, or groupy.

We proceed by discussing each type in turn, from the most prevalent to the least prevalent. The most prevalent type—estimated 22.7% of subjects—places slightly higher

33. In the following discussion, we describe the estimated fractions of subjects generated by the mixing model. Section VI considers the subjects assigned to each type by their ex post probabilities.

34. The online Appendix reports (1) the predicted difference in own income when facing an ingroup vs. an outgroup participant, and (2) the absolute payoff levels that underly these ingroup/outgroup differences; i.e., the predicted payoffs for each type in each match, the predicted payoffs for an ingroup recipient and an outgroup recipient facing each type, as well as predicted expected total surplus for each type. All these payoffs generally track the utility function parameters. The tests of the cross-type differences in all these amounts are also reported in the online Appendix; these tests have low power for types 6,7,8, and 9 due to small numbers.

weight on others' income than own income, when the recipient is ingroup or the recipient is outgroup, in both group treatments.³⁵ In each group treatment, the predicted favoritism levels are not statistically different than zero, and we cannot reject that the favoritism levels are the same at the 5% level. By the criteria, Type 1 is then not groupy.

Type 2, the next largest type—estimated 19.6% of subjects—favors strongly own income relative to others' income, ingroup or outgroup, in both group treatments.³⁶ The political group utility parameters are statistically different than those for the minimal group treatment; political group predicted favoritism is statistically positive while minimal group favoritism is statistically zero. This difference is statistically different than zero, hence this type is conditionally groupy. However, this label is somewhat misleading. The difference in the favoritism levels is small and arguably not economically significant, at only 4.4% of the total possible favoritism. Indeed, the difference is just below that of Type 1. If we impose an economic significance criterion of the difference in favoritisms to be at least 5%, this type might be considered groupy. In some of the comparisons below, we will consider two possible interpretations of Type 2, as either conditionally groupy or groupy.³⁷

35. These weights are consistent with achieving equal levels of income between self and other subjects. Predicted payoffs own for a decision-maker of this type are the lowest of all types, though predicted aggregated payoffs are the highest among the nine types (see the online Appendix).

36. Type 2 earns the highest payoffs of all nine types, as shown in Tables OA4 and OA5 of the online Appendix.

37. At the same time, the groupy designation is also somewhat misleading, since favoritism levels are very small, and well below the favoritism levels of other groupy types (discussed below). With consistently high weights on own income relative to others' incomes, this type's own income is statistically the same when allocating to ingroup and outgroup (as reported in Table OA3 in the online appendix). This type is then possibly a case where including own income in the groupiness definition could change its groupiness designation.

Like Type 2, the third most prevalent type—estimated 17.2% of subjects—places higher weight on own income than other’s income, but the magnitudes are smaller. Relative to the minimal group, Type 3’s weight in the political group on ingroup income is higher and weight on outgroup income is lower. The predicted favoritism level in political group is both economically and statistically significant. With minimal group favoritism statistically zero, this type is unambiguously conditionally groupy.

The fourth most prevalent type—estimated 12.1% of subjects—weighs ingroup recipient’s income much more than outgroup recipient’s income, and the utility parameters are statistically the same in the minimal and political group treatment. Consequently, this type has the largest favoritism levels in both minimal group and political group treatment, up to 60% of the maximal difference, and this type is unambiguously groupy.

The fifth type—estimated 10.7% of subjects—is also groupy, but mildly so; the utility parameters are statistically different across group conditions, but translate into positive but small favoritism, about 10% of the total possible, and is statistically the same in the two group conditions.

The sixth type—estimated 8.5% of subjects—is similar to Type 3, but with larger divergence between parameters in minimal group and political group vis-à-vis outgroup recipient. This type, with a difference in favoritism that is the largest among the types at 24.2%, is unambiguously conditionally groupy.

For the seventh (4.3% of subjects), eighth (3.6% of subjects), and ninth (1.4% of subjects) types, statistical tests have low power, yet we can still see the distinctive behavior of each. Type 7 is the only type whose political vs. minimal group utility

parameters only differ for own income and vis-à-vis ingroup; this difference though leads a small difference in favoritism, and by the criteria, this type is not groupy.

Type 8, unlike all other types, puts statistically significant negative weight on recipient's income, whether ingroup, outgroup, political group or minimal group. With the small weight on own income, this type appears to be willing to sacrifice own income in order to reduce the income of others, and more so in the political group.³⁸ The resulting favoritism difference, however, is not statistically significant, so this type is groupy.

Finally, Type 9 appears to care only about own income when giving to an ingroup participant and to care about own and other's income when giving to an outgroup participant, and these weights are higher in the political than the minimal group; by the criteria this type is not groupy.

For the average subject, neither the utility function nor the favoritism levels match those of any one of the nine types. The favoritism level in minimal group condition is closest in magnitude to that of Type 5 but political group favoritism is closest to that of Type 3 (and for each of these levels, we cannot reject that the respective magnitudes are the same).³⁹ The average subject is unambiguously conditionally groupy, but estimated less than half (45%) of subjects are conditionally groupy if Type 2 is included, and only 26% if Type 2 is not included.

38. Matches involving Type 8 yield the lowest aggregate payoffs, as reported in the online Appendix.

39. Tests for differences for all the cross-type comparisons are in the online Appendix. The average subject's minimal group favoritism level is statistically different than those of Types 1, 2, 3, and 4 at the 0.05 level or below; political group favoritism is statistically different than those of Types 1, 2, 4, and 5 at the 0.05 level or below. Many of the tests relative to Types 6,7,8, and 9 are not sufficiently well powered to detect differences.

Figure 5, below, shows the graph of predicted favoritism for the nine types and for the average subject. The average subject's predicted favoritism levels are marked with an "A." Circles indicate the type is not groupy, diamonds indicate the type is conditionally groupy, and squares indicate the type is groupy. The dashed line is the 45⁰ line. Summing over the estimated proportions of subjects of each type, 29% are not groupy, 45% are conditionally groupy, and 26% are groupy.⁴⁰ These estimated proportions show the greater precision accomplished by the latent class model relative to tests on subjects' individual data reported in the previous section.

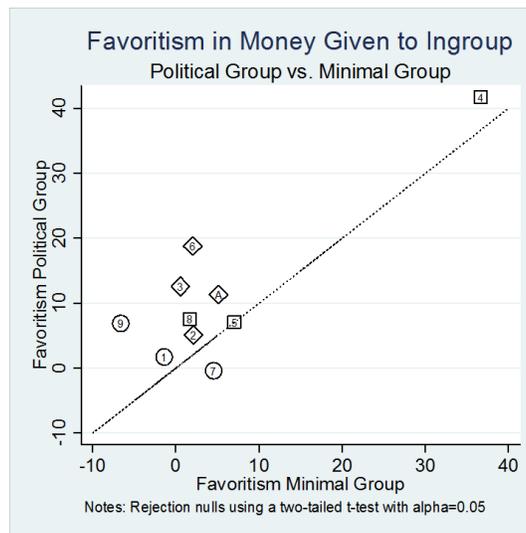


Figure 5.
Predicted Favoritism Levels for the Nine Types and the Average

Before turning to the demographics of each type, we study possible order effects, which can arise in any within subject experiment. In Table 3, we compare the prevalence of types among those who received the minimal group treatment first with the prevalence of types among those who received the political group treatment first. The comparison

40. The online Appendix provides the cross tabulation between subjects' type assignments and subjects' groupiness designations from their individual data.

shows suggestive evidence of a spillover of the more salient, political treatment on the minimal group treatment for one set of subjects. The fraction of subjects of each type are statistically identical except for Type 2, which has fewer political-group-first subjects, and Type 4, which has more political-group-first subjects. Subjects who have the highest levels of favoritism in the political treatment are more likely to have the highest levels of favoritism in the minimal group treatment if they received the political treatment first. The order of group treatments does not appear to matter for the prevalence of any other types.⁴¹

VI. Demographic and Behavior Correlates of Utility Types

By classifying subjects into types, we can study demographic and other correlates of individual variation in behavior in the group treatments.⁴² We exploit three sources of individual information. The first source is the data from the questionnaire administered at the end of the experiment, which asked standard demographic questions as well as queried subjects on some personal practices and attitudes. The second source is the data from the survey administered at the beginning of the political treatment, which asked subjects' political party affiliations and political positions. The third source is the data from the non-group condition, which occurred at the start of the experiment for all subjects, where subjects allocated income to themselves and to an anonymous randomly selected other participant in the experiment.

41. As another check, we split the sample into minimal-group-first subjects and political-group-first subjects and estimate separately a nine-type latent class model for each. The predicted behaviors of the most prevalent types are qualitatively the same for both subsamples; estimates for the less prevalent types are similar but harder to compare due to small numbers. The online appendix provides the predicted favoritism levels from these estimations.

42. Note that, due to the discreteness of assigning subjects, the fraction of subjects assigned to each type do not exactly match the estimated proportions from latent class model reported above.

We begin with the first two sources and focus on individual demographics as well as out-of-laboratory choices that could relate to groupiness. Groupiness could vary by sex, ethnicity, or family income. Group-oriented practices and attitudes, such as religious service attendance, political party affiliation, and distrust of strangers, could be correlated with behavior in the lab. Table 4 reports on the composition of each of the nine types in terms of gender, race, birthplace (US or not), parents' education, political party affiliation (Democrat, Republican, Independent), distrust of strangers, and frequency of religious service attendance. We also combine the types according to their groupiness designations and conduct the same tests, with results reported in Table 5. Looking across these tables, we see the types are largely similar in terms of characteristics, with perhaps two exceptions related to the groupy types. Groupy types consist of fewer women (49%) than conditionally groupy types (75%). The most biased groupy type – Type 4 – has fewer subjects with fathers with advanced degrees (13%) than Type 2 – the weakest conditionally groupy type (50%).⁴³

With little of the individual demographics or outside-the-laboratory choices providing leverage on groupiness, we ask whether individual behavior in the non-group condition is informative of behavioral patterns in the group treatments. To do so, we estimate individual social preferences in the non-group condition: we posit the same simple utility function as above, estimate a series of latent class models, for which the selection criteria suggest six types, and categorize subjects into types by the posterior probabilities. To better see any relationship between subjects' social preferences in the

43. The methodology in Kranton, Pease, Sanders, and Huettel (2016) and the definition of groupy is substantively different than in the present study. The finding of the relationship between not groupy and political independence does not appear in this analysis which divides the subject pool into smaller sets.

non-group condition and their groupiness in the group treatments, we merge the six types in the non-group condition into three sets according to the utility function parameter values.⁴⁴ For the first set, the weight on own payoffs is higher than weight on other's payoffs (62% of subjects); for the second set, weight on own payoffs is lower than the weight on others' payoffs (33% of subjects); for the third set, weight on others' payoffs is negative (5% of subjects). We then cross-tabulate these three sets of subjects with their type in the group treatments.

Figure 6 gives a bar graph of the tabulations, where the Types 1–9 are placed according to their groupiness designation, then by prevalence. The Figure shows that social preferences in the non-group condition are limitedly indicative of groupiness. Subjects who put negative weight on other's income in the non-group condition, whom we might call “mean,” are all groupy in the group conditions and most are Type 8. Such social preferences could relate, then, to a sort of groupiness.⁴⁵ On the other hand, subjects who put balanced weights on own and other's income in the non-group condition, are spread between not groupy and conditionally groupy types. Some of these subjects could have preference for fairness that carries into both group conditions, but a substantial subset is Type 6, which has the largest bias against the political outgroup. Subjects who are “selfish” in the non-group condition, putting large weight on own

44. The full estimation of the non-group treatment utility functions is provided in the on-line appendix along with the six-type by nine-type cross-tabulation.

45. Several researchers have identified similar “mean” social preferences in non-group settings. The behavior has been given different names, e.g., “spitefulness,” “competitiveness,” “nastiness,” and “equity aversion” (e.g., Levine (1998), Fehr, Hoff, and Kshetramde (2008)), Abbink & Sadrieh (2009), Ibierrri & Rey-Biel (2012), Fershtman, Gneezy, & List (2012)). Kranton, Pease, Sanders, and Huettel (2016), which estimates a utility function with weights on own income and on relative income, finds that in the group conditions about 20% of subjects give up own income in order to choose allocations with greater differences between their income and others' incomes.

income, are spread mostly between conditionally groupy and groupy types. Overall, this exercise indicates social preferences per se are a separate phenomenon from groupiness.⁴⁶

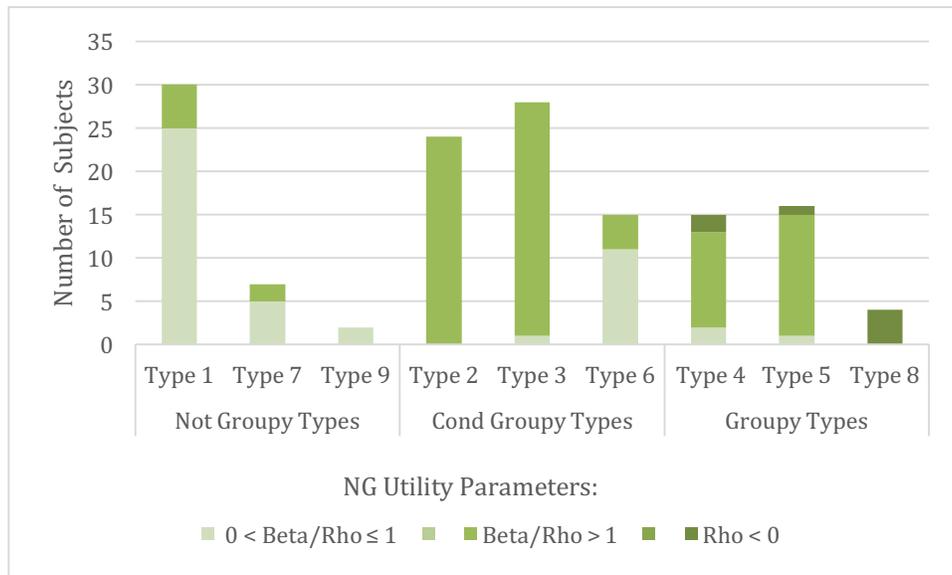


Figure 6.
Non-Group Social Preferences and Groupiness

VII. Conclusion

From the sandlot, where a friendly pick-up game can turn into a brawl, to the public square where a democracy movement can turn into a civil war, people in groups alternatively coalesce or conflict. This experiment studies individual behavior in different given group settings. The study builds on the long tradition of experiments in social psychology on group conflict and on the established literature in economics on income allocation and social preferences. The experiment strips away social interactions,

46. We also estimated a seemingly unrelated regression where the posterior probability of being each of the nine type was modeled as a function of the posterior probability of having one of three sets of specified non-group utility function parameters. (The six estimated types were merged into three sets, as in Figure 6.) The results, reported in the online Appendix, confirm statistically the interpretation of the cross tabulation and its graph reported here.

punishments, collective benefits and other dynamics that might drive people to help or hurt others in different groups. The simplicity of the task places the focus on individuals' underlying predispositions. With a new design and methods, the paper asks whether individuals might be more or less prone to treat people differently depending on others' group affiliation.

In the experiment, subjects choose allocations of income to self and others. Each subject allocates income in a non-group condition and then two group treatments—minimal group and political group—in both ingroup pairings and outgroup pairings. We study the behavior across the two group treatments, and in particular, the differences in amounts of income given to ingroup and outgroup subjects. Estimating a series of latent class models, we identify nine distinct behavioral patterns and classify individual subjects as these behavioral types.

The results reveal systematic divergent responses to group treatments. Most subjects are either groupy—always favoring the ingroup—or not groupy—never favoring the ingroup. While all not groupy subjects are alike in that allocations do not differ by group, they can be distinct in that they make consistently different allocations – with more or less income given to other subjects. The groupy subjects are all alike in that favor the ingroup, but the levels of favoritism have a wide range. The remaining set of subjects is conditionally groupy, biasing allocations only in the political treatment, again with a range.

Conducted in a lab at a university, the group divisions in this experiment, including the political groups, are obviously pallid compared to the political and ethnic divisions that rend societies. More salient and immediate group divisions could of

course lead people to favor their ingroup. Nonetheless, the results of the experiment call for a richer model of bias—one that includes individual predilections as key variables. Outside the laboratory, people also respond in divergent ways to leaders and policies that promote and foment group divisions; some people join organizations that actively harm outgroups, some shelter outgroups at their own peril, and some people simply seek to profit in all events. The range of behavior of subjects in this experiment mirrors this diversity. Future experimental research could test whether individual groupiness is a robust trait, both across time and across tasks, such as the provision of public goods or exchanges enhanced by trust. Conducting experiments both inside and outside the laboratory and empirical studies could further investigate demographic, psychometric, and cultural correlates of groupiness.⁴⁷

47. In an MTurk study, Kranton and Sanders (2017) finds that no “big five” personality measure relates to biased allocations in a minimal group setting, and hence would be unrelated to groupiness.

References

- Abbink, Klaus and Abdolkarim Sadrieh. 2009. "The pleasure of being nasty," *Economics Letters* 105(3), pp. 306-308.
- Akerlof, George and Rachel Kranton. 2000. "Economics and Identity," *The Quarterly Journal of Economics* CVX (3), August, pp. 715-753.
- Akerlof, George and Rachel Kranton. 2010. *Identity Economics*, Princeton: Princeton University Press: 2010.
- Andreoni, James and John Miller. 2002. "Giving according to GARP: An Experimental Test of the Consistence of Preferences for Altruism," *Econometrica* 70(2), pp. 737-753.
- Bernhard, Helen, Ernst Fehr, and Urs Fischbacher. 2006. "Group Affiliation and Altruistic Norm Enforcement," *American Economic Review*, 96(2), pp. 217-21.
- Berreby, David. 2008. *Us and Them: The Science of Identity*, Chicago: University of Chicago Press.
- Bolton, Gary & Axel Ockenfels. 2000. "ERC: A Theory of Equity, Reciprocity, and Competition," *American Economic Review* 90 (1), pp. 166-193.
- Bosch-Domènech, Antoni, José García-Montalvo, Rosemarie Nagel and Albert Satorra. 2010. "A Finite Mixture Analysis of Beauty-Contest Data Using Generalized Beta Distributions," *Experimental Economics*, December 13(4), pp. 461-475.
- Bruhin, Adrian, Lorenz Goette, Adrian Roethlisberger, Alexander Markovic, Regula Buchli, and Beat M. Frey. 2015. "Call of duty: the effects of phone calls on blood donor motivation," *Transfusion*, 55, pp. 2645-2652.
- Bruhin, Adrian, Ernst Fehr, and Daniel Schunk. 2016. "The Many Faces of Human Sociality: Uncovering the Distribution and Stability of Social Preferences," *CEifo Working Paper Series No 5744*.
- Burnham, Kenneth P and David R. Anderson. 2004. "Multimodel Inference: Understanding AIC and BIC in Model Selection," *Sociological Methods & Research*, 33, pp. 261-304.
- Camerer, Colin F., and Ernst Fehr. 2004 "Measuring Social Norms and Preferences Using Experimental Games: A Guide for Social Scientists." In *Foundations of Human Sociality: Economic Experiments and Ethnographic Evidence from Fifteen Small-Scale Societies*, eds. Joseph Henrich et al., 55–95. Oxford: Oxford University Press.

- Charness, Gary and Matthew Rabin. 2002. "Understanding Social Preferences with Simple Tests," *The Quarterly Journal of Economics* 117(3), August 2002, pp. 817-869.
- Charness, Gary, Luca Rigotti, and Aldo Rustichini. 2007. "Individual Behavior and Group Membership," *American Economic Review* 97(4), pp. 1340-1352.
- Chen, Roy and Yan Chen. 2011. "The Potential of Social Identity for Equilibrium Selection." *American Economic Review* 101 (6) October, pp. 2562--2589.
- Chen, Yan and Sherry Li. 2009. "Group Identity and Social Preferences," *American Economic Review* 99(1), pp. 431-457.
- Conte, Anna & Hey, John D. & Moffatt, Peter G., 2011. "Mixture models of choice under risk," *Journal of Econometrics*, 162(1), May, pp. 79-88.
- Engelmann, Dirk and Martin Strobel. 2004. "Inequality Aversion, Efficiency, and Maximin Preferences in Simple Distribution Experiments," *American Economic Review* 94 (4), pp. 857-869.
- Fehr, Ernst and Klaus M. Schmidt. 1999. "A Theory of Fairness, Competition, and Cooperation," *Quarterly Journal of Economics* 114(3), pp. 817-68.
- Fehr, Ernst, Karla Hoff, and Mayuresh Kshetramade. 2008. "Spite and Development," *American Economic Review: Papers & Proceedings* 98(2), 494-499.
- Fershtman, Chaim & Uri Gneezy. 2000. "Discrimination in a Segmented Society: An Experimental Approach," *Quarterly Journal of Economics*, 116(1), February, 351-377.
- Fershtman, Chaim, Uri Gneezy, and John List. 2012. "Equity Aversion: Social Norms and the Desire to Be Ahead." *AEJ: Microeconomics*, 4(4), 131-44.
- Fischbacher, Urs, Ralph Hertwig, and Adrian Bruhin. 2013. "How to Model Heterogeneity in Costly Punishment: Insights from Responders' Response Times," *Journal of Behavioral Decision Making*, 26, pp.462-476.
- Fisman, Raymond, Shachar Kariv, and Daniel Markovits. 2007. "Individual Preferences for Giving." *American Economic Review*, 97(5), pp.1858-1876.
- Fowler, James H. and Cindy D. Khan. 2007. "Beyond the Self: Social Identity, Altruism, and Political Participation," *The Journal of Politics*, 69(3), pp. 813-827.
- Glaeser, Edward and David Laibson, Jose Scheinkman, Christine Souter. 2000. "Measuring Trust," *Quarterly Journal of Economics*, 115(3), pp. 811-846.

- Goette, Lorenz, David Huffman, and Stephan Meier. 2006. "The Impact of Group Membership on Cooperation and Norm Enforcement: Evidence Using Random Assignment to Real Social Groups." *American Economic Review*, 96(2), pp. 212-16.
- Goette, Lorenz, David Huffman, and Stephan Meier. 2012. "The Impact of Social Ties on Group Interaction: Evidence from Minimal Groups and Randomly Assigned Real Groups." *American Economic Journal: Microeconomics*, 41(1), pp. 101-115.
- Harrison, Glenn and Elizabeth Rutström. 2009 "Expected Utility And Prospect Theory: One Wedding and Decent Funeral," *Experimental Economics* 12(2), June 2009, 133-158.
- Hargreaves Heap, Shaun P., and Daniel John Zizzo. 2009. "The Value of Groups." *American Economic Review*, 99(1), pp. 295-323.
- Hornsey, Matthew J. 2008. "Social Identity Theory and Self-Categorization Theory: A Historical Review," *Social and Personality Psychology Compass* 2/1, pp. 204-222.
- Iriberry, Nagore and Pedro Rey-Biel. 2013. "Elicited Beliefs and Social Information in Modified Dictator Games," *Quantitative Economics*, 4, pp. 515–547.
- Klor, Esteban and Moses Shayo. 2010. "Social Identity and Preferences over Redistribution," *Journal of Public Economics* 94(3-4), pp. 269-278.
- Kranton, Rachel, Matthew Pease, Seth Sanders, and Scott Huettel. 2013. "Identity, Groups and Social Preferences," working paper, <https://sites.duke.edu/rachelkranton/>
- Kranton, Rachel, Matthew Pease, Seth Sanders, and Scott Huettel. 2016. "Groupy and Not groupy Behavior: Deconstructing Bias in Social Preferences," working paper, <https://sites.duke.edu/rachelkranton/>.
- Kranton, Rachel and Seth Sanders. 2017. "Groupy vs. Non-Groupy Behavior: Personality, Region, and Party Affiliation," *American Economic Review Papers and Proceedings*, Vol 107(5), May, pp. 65-69
- Levine, David K. 1998. "Modeling Altruism and Spitefulness in Experiments," *Review of Economic Dynamics* 1, pp. 593-622.
- Levitt, Steven and John List. 2007. "What do Laboratory Experiments Measuring Social Preferences Reveal about the Real World?," *Journal of Economic Perspectives*. Vol 21(2), Spring, pp. 153-174.

- Pechar, Emily and Rachel Kranton. 2017. "Moderators of Intergroup Discrimination in the Minimal Group Paradigm: A Meta-Analysis," working paper, Duke University.
- Nagin, Daniel. 2005. *Group-Based Modeling of Development*. Cambridge: Harvard University Press.
- Sherif, Muzafer, O.J. Harvey, B. Jack White, William R. Hood, and Carolyn W. Sherif. 1961. "The Robbers Cave Experiment: Intergroup Conflict and Cooperation." Wesleyan University Press:
- Stahl, D. O. 1996. "Boundedly Rational Rule Learning in a Guessing Game," *Games and Economic Behavior* 16, pp. 303-330.
- Stahl, Dale and Paul Wilson. 1995. "On Players' Models of Other Players: Theory and Experimental Evidence," *Games and Economic Behavior* 10, pp. 218-254.
- Tajfel, Henri, M. G. Billig, R. P. Bunding, and Claude Flament. 1971. "Social categorization and intergroup behaviour," *European Journal of Social Psychology*, Vol 1 (2), April/June, pp. 149-178.
- Tajfel, Henri and John Turner, "An Integrative Theory of Intergroup Conflict," in Stephen Worchel and William Austin, eds., *The Social Psychology of Intergroup Relations*, Monterey, CA: Brooks/Cole, 1979.

Table 1. Results of Latent Class Model with Nine Types: Utility Parameters and Proportions of the Subject Pool

Utility Function Parameters	Type 1	Type 2	Type 3	Type 4	Type 5	Type 6	Type 7	Type 8	Type 9	Average Subject
Beta (MG, In)	0.00917*** (0.00250)	0.213*** (0.0190)	0.0706*** (0.00652)	0.0112*** (0.00284)	0.0728*** (0.00784)	0.00530 (0.00335)	0.0163*** (0.00495)	0.0813*** (0.0158)	0.0312*** (0.0108)	0.0282*** (0.00124)
Rho (MG, In)	0.0229*** (0.00155)	0.0209*** (0.00299)	0.0188*** (0.00211)	0.0106*** (0.00148)	0.00223 (0.00198)	0.0133*** (0.00186)	-0.00197 (0.00224)	-0.0493*** (0.00932)	-0.00665 (0.00420)	0.00938*** (0.000553)
Beta (MG, Out)	0.0112*** (0.00274)	0.231*** (0.0273)	0.0597*** (0.00538)	0.0369*** (0.00455)	0.0735*** (0.00792)	0.0114*** (0.00367)	0.00489 (0.00465)	0.0536*** (0.0116)	0.0228** (0.00921)	0.0270*** (0.00121)
Rho (MG, Out)	0.0259*** (0.00168)	0.0148*** (0.00273)	0.0160*** (0.00182)	-0.0264*** (0.00242)	-0.00704*** (0.00213)	0.0122*** (0.00182)	-0.00572** (0.00230)	-0.0433*** (0.00740)	-0.000872 (0.00408)	0.00460*** (0.000496)
Beta (MG, In)-	0.00106 (0.00360)	-0.00518 (0.0262)	-0.0222*** (0.00770)	-0.00317 (0.00406)	0.0629*** (0.0166)	0.00838* (0.00480)	-0.0184*** (0.00664)	-0.0387** (0.0179)	-0.0530*** (0.0170)	-0.00222 (0.00171)
Rho (MG, In) -	0.00124 (0.00220)	0.00101 (0.00426)	0.000584 (0.00278)	0.00319 (0.00216)	-0.00797** (0.00318)	-0.00278 (0.00260)	-0.00182 (0.00313)	0.0304*** (0.0100)	0.0409*** (0.00940)	0.00104 (0.000777)
Beta (MG, Out) -	0.00111 (0.00370)	-0.0878*** (0.0294)	0.0281*** (0.00894)	0.00732 (0.00681)	0.0354** (0.0149)	0.0118** (0.00535)	-0.000522 (0.00644)	0.00103 (0.0161)	-0.0386*** (0.0137)	0.00397*** (0.00177)
Rho (MG, Out)-	-0.00406* (0.00227)	-0.0121*** (0.00324)	-0.00874*** (0.00244)	-0.00475 (0.00373)	-0.0109*** (0.00335)	-0.0161*** (0.00244)	0.00263 (0.00318)	0.00473 (0.00990)	0.0198*** (0.00674)	-0.00388*** (0.000698)
Fraction	0.227*** (0.0354)	0.196*** (0.0344)	0.172*** (0.0323)	0.121*** (0.0274)	0.107*** (0.0269)	0.0852*** (0.0237)	0.0425** (0.0170)	0.0355** (0.0156)	0.0142 (0.00996)	1 -
Observations	14,587	14,587	14,587	14,587	14,587	14,587	14,587	14,587	14,587	14,587

Standard errors in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

Table 2: Predicted Favoritism Levels for each Type and Groupness Designations

Group Treatment/Designation	Type 1	Type 2	Type 3	Type 4	Type 5	Type 6	Type 7	Type 8	Type 9	Average Subject
MG	-1.415 (1.191)	2.096** (0.955)	0.555 (1.521)	36.65*** (1.918)	7.020*** (2.071)	2.017 (2.618)	4.513 (4.221)	1.683 (1.985)	-6.642 (6.717)	5.137*** (0.788)
POL	1.692 (1.242)	5.129*** (1.143)	12.57*** (1.503)	41.71*** (1.774)	7.082*** (1.618)	18.78*** (2.778)	-0.399 (4.302)	7.547*** (2.729)	6.892 (4.438)	11.31*** (0.773)
POL minus MG: Absolute	3.107* (1.740)	3.033** (1.407)	12.02*** (2.091)	5.055* (2.612)	0.0617 (2.638)	16.76*** (3.820)	-4.911 (6.019)	5.863* (3.374)	13.53* (8.051)	6.173*** (1.104)
POL-MG Percent	4.5%	4.4%	17.4%	7.3%	0.1%	24.2%	7.1%	8.5%	19.5%	8.9%
Groupness - Statistical Test	Not Groupy	Cond Groupy	Cond Groupy	Groupy	Groupy	Cond Groupy	Not Groupy	Groupy	Not Groupy	Cond Groupy
Groupness - Econ Significance	Not Groupy	Groupy	Cond Groupy	Groupy	Groupy	Cond Groupy	Not Groupy	Groupy	Not Groupy	Cond Groupy
Fraction of Subjects	0.227	0.196	0.172	0.121	0.107	0.085	0.043	0.036	0.014	1

Standard errors in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

Table 3. Prevalence of Subjects of each Type by Order of Group Treatment

Order	Type 1	Type 2	Type 3	Type 4	Type 5	Type 6	Type 7	Type 8	Type 9
POL FIRST Fraction	0.193*** (0.0522)	0.103** (0.0406)	0.140*** (0.0464)	0.210*** (0.0538)	0.175*** (0.0508)	0.123*** (0.0438)	0.0392* (0.0237)	0.0175 (0.0173)	
MG FIRST Fraction	0.253*** (0.0478)	0.270*** (0.0506)	0.184*** (0.0431)	0.0602** (0.0261)	0.0640** (0.0292)	0.0606** (0.0263)	0.0452** (0.0217)	0.0482** (0.0235)	0.0151 (0.0113)
Difference in Fraction	-0.0600 (0.0708)	-0.167*** (0.0647)	-0.0441 (0.0633)	0.149** (0.0598)	0.111* (0.0583)	0.0622 (0.0511)	-0.00597 (0.0298)	-0.0307 (0.0292)	
Fraction of Subject Pool	0.227	0.196	0.172	0.121	0.107	0.085	0.043	0.036	0.014

Standard errors in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

Table 4. Demographics and Nine Types

Fraction of Type	Non-Groupy			Cond Groupy			Groupy			Total Pool
	Type 1	Type 7	Type 9	Type 2	Type 3	Type 6	Type 4	Type 5	Type 8	
Female	0.60	0.71	0.50	0.75	0.71	0.80	0.60	0.44 ⁽⁶⁾	0.25 ^(2,6)	0.65
African American	0.17	0.57	0.00 ^(1,4,6,7)	0.04 ^(6,7)	0.07 ⁽⁷⁾	0.33	0.27	0.19	0.50	0.18
Born in United States	0.86	1.00 ^(1,2,3,5)	0.50	0.83	0.79	0.87	0.87	0.67	0.75	0.82
Mostly Distrust Strangers	0.45	0.29	0.00 ^(1,2,3,4,6)	0.38	0.18 ⁽¹⁾	0.27	0.47	0.13 ^(1,4)	0.00 ^(1,2,3,4,6)	0.30
No Religious Attendance	0.07	0.14	0.00	0.04	0.04	0.20	0.13	0.07	0.00	0.08
Democrat	0.40	0.57	0.00 ^(1,2,3,4,5,6,7,8)	0.58	0.57	0.47	0.53	0.25 ^(2,3)	0.75	0.48
Independent	0.50	0.14 ^(1,5)	0.50	0.38	0.21 ^(1,5)	0.47	0.40	0.56	0.25	0.39
Lived with both parents	0.79	0.57	1.00 ^(1,2,3,4,6,7)	0.83	0.75	0.73	0.73	0.87	1.00 ^(1,2,3,4,6,7)	0.78
Father advanced degree	0.30	0.00 ^(1,2,3,5,6)	0.50	0.50	0.32	0.33	0.13 ⁽²⁾	0.50	0.00 ^(1,2,3,5,6)	0.33
Mother advanced degree	0.17	0.14	0.50	0.29	0.07 ⁽²⁾	0.20	0.00 ^(1,2,5)	0.25	0.00 ^(1,2)	0.16

Superscripts denote fraction is significantly different at the 0.05 level from noted types using a t-test.

Table 5. Demographics and Groupiness

Fraction of Type	Not Groupy	Cond Groupy	Groupy	All Subjects
Female	0.62	0.75	0.49 ^(CG)	0.65
African American	0.23	0.12	0.26	0.18
Born in United States	0.87	0.82	0.77	0.82
Mostly Distrust Strangers	0.40	0.27	0.27	0.30
No Religious Attendance	0.08	0.07	0.09	0.08
Democrat	0.41	0.55	0.43	0.48
Independent	0.44	0.33	0.46	0.39
Lived with both parents	0.76	0.78	0.82	0.78
Father advanced degree	0.26	0.39	0.29	0.33
Mother advanced degree	0.18	0.18	0.11	0.16

Superscripts denote fraction is significantly different at the 0.05 level from noted types using a t-test.

Appendix

Instructions to Participants

PAGE 1

WELCOME!

INSTRUCTIONS

Thank you for participating in this experiment. The object of this investigation is to study how people make decisions. There is no deception in this experiment – and we want you to understand everything about the procedures. If you have any questions at any time, please ask the experiment organizer in the room.

PART I: THE CHOICE TASK

A) During the experiment, you will be presented with a series of choices. For each choice, you will be asked to award points to between either (1) yourself and another participant or (2) two other participants. You will earn the points you allocate to yourself, and the other person will earn the points you allocate to him or her. At the end of the experiment, one of your choices will be selected at random by a computer and the points earned will be converted into payments.

Each decision is independent from the others. Your decisions and outcomes in one choice will not affect your outcomes in any other choice. For each choice, you will be paired with new participants.

Use LEFT and RIGHT arrow keys to make your choices.

PART II and III:

A) INITIAL SURVEY

You will take a brief survey. There are no right or wrong answers. Your answers to these questions will not affect your payments. Please only use the RIGHT and LEFT arrow keys or NUMBER keys as instructed to answer all questions.

B) THE CHOICE TASK

After completing the initial survey, you will once again be presented with a series of choices. You will be anonymously paired with two new participants. These participants will remain the same throughout this part of the experiment. At the end of the experiment, one of your choices will be selected at random by a computer and the points earned will be converted into payments. Each decision is independent from the others. Your decisions and outcomes in one choice will not affect your outcomes in any other choice.

TURN PAGE OVER FOR ADDITIONAL INSTRUCTIONS

PAGE 2

PAYMENT

At the end of the experiment, the points you get will be converted into money by a predetermined conversion factor. This money will be added to your \$6 participation payment and given to you at the end of the experiment. Since we want you to focus on completing the experiment and not calculating points to money conversions, we will not inform you of the conversion factor. However, we expect participants to earn between \$12 and \$18, with an average of \$15.

SETUP

You will make all choices on a computer screen. You will make approximately 200 choices.

For each choice, you will see a screen that presents the two different points allocations you can make.

	YOU	OTHER
GREEN	10	10
BLUE	15	5

After a one second pause, two arrows will appear so you can pick which allocation you prefer. You can press either 'LEFT' or 'RIGHT' arrow keys on the keyboard to match the arrows presented on the screen. Please only touch the RIGHT or LEFT arrow keys for all choices.

	YOU	OTHER
GREEN	10	10
BLUE	15	5
	←Green	Blue→

Are there any questions? Press any key to begin.

Table A1. Normalized Matrices

A. Matrices where $\pi_i \geq \pi_j$ in both rows, ordered by $\Delta\pi_i/(\Delta\pi_i-\Delta\pi_j)$

Matrix Number	(π_i, π_j) (π'_i, π'_j)	$\Delta\pi_i/(\Delta\pi_i-\Delta\pi_j)$
14	140 100	-2
	100 40	
12	140 100	-1.5
	80 0	
16	140 100	-0.5
	120 40	
19	140 140	-0.5
	120 80	
15	140 100	-0.25
	120 0	
18	140 140	-0.16
	120 0	
1	100 100	0
	100 20	
7	140 20	0.2
	120 100	
9	140 40	0.2
	120 120	
10	140 60	0.33
	120 100	
11	140 80	0.33
	120 120	
21	160 0	0.375
	100 100	
5	120 80	0.5
	100 100	
22	160 40	0.5
	120 80	
25	200 0	0.5
	100 100	
26	200 0	0.5
	180 20	
8	140 40	0.6
	80 80	
17	140 120	3
	80 80	
13	140 100	NA
	80 40	

B. Matrices where $\pi_i < \pi_j$ in at least one rows ordered by $\Delta\pi_i/(\Delta\pi_i-\Delta\pi_j)$

Matrix Number	(π_i, π_j) (π'_i, π'_j)	$\Delta\pi_i/(\Delta\pi_i-\Delta\pi_j)$
---------------	--	---

3	100	200	0
	100	100	
4	100	200	0
	100	140	
2	100	140	0
	100	60	
6	140	0	0.125
	120	140	
23	160	80	0.2
	140	160	
20	140	140	0.33
	120	180	
24	160	120	0.33
	140	160	

Table A.2. Distribution of Subjects' Political Affiliations and Leanings

SURVEY CATEGORY	% OF SUBJECTS
Democrat – Strong	15
Democrat – Moderate	33
Republican – Strong	0
Republican – Moderate	13
Independent – Dem leaning	13
Independent – Rep leaning	10
None of the Above – Dem leaning	11
None of the Above – Rep leaning	5

Table A3. Bayesian Inference Criterion Calculations for Latent Class Models

Features of Latent Class Model	7	8	9	10	11
Number of types	7	8	9	10	11
Number of Parameters	63	72	81	90	99
Log Likelihood	-6194.96	-6055.95	-6015.03	-6003.85	-5980.98
BIC	12993.95	12802.23	12806.69	12870.61	12911.17
Observations	14587	14587	14587	14587	14587