

Distributional Preferences and Political Behavior*

Raymond Fisman, Pamela Jakiela, and Shachar Kariv[†]

June 25, 2016

Abstract

We decompose distributional preferences into fair-mindedness (weight on oneself vs. others) and equality-efficiency tradeoffs, and measure both at the individual level in a large and diverse sample of Americans. We find considerable heterogeneity in both the extent of fair-mindedness and willingness to trade off equality and efficiency, much of which cannot be explained by standard socioeconomic or demographic variables. Subjects' equality-efficiency tradeoffs predict their political decisions: equality-focused subjects are more likely to have voted for Barack Obama, and to be affiliated with the Democratic Party. Our findings shed light on how American voters are motivated by their distributional preferences.

JEL Classification Numbers: C91, D64.

Keywords: Distributional preferences, social preferences, fair-mindedness, self-interest impartiality, equality, efficiency, redistribution, political decisions, voting, party affiliation, American Life Panel (ALP), experiment

*Previously circulated as “The Distributional Preferences of Americans.” We thank Daniel Markovits, Daniel Silverman, and conference participants at Lund University and the Stanford Institute for Theoretical Economics for helpful discussions and comments. We thank the American Life Panel (ALP) team at the RAND Corporation for software development and technical and administrative support. We acknowledge financial support from the Center for Equitable Growth (CEG) at the University of California, Berkeley.

[†]Fisman: Columbia University (email: rfisman@bu.edu); Jakiela: University of Maryland (email: pjakiela@umd.edu); Kariv: University of California, Berkeley (email: kariv@berkeley.edu).

1 Introduction

Distributional preferences shape individual opinions on a range of issues related to the redistribution of income — examples include social security, unemployment benefits, and government-sponsored healthcare. These issues are complex and contentious in part because people promote their competing private interests, but they also often disagree about what constitutes a just or equitable outcome, either in general or in particular situations. We therefore cannot understand public opinion on a number of important policy issues without understanding the individual distributional preferences of the general population.

Distributional preferences may naturally be divided into two qualitatively different components: the weight on own income versus the incomes of others (fair-mindedness), and the weight on reducing differences in incomes (equality) versus increasing total income (efficiency). Political debates often center on the redistribution of income, and fair-minded people may disagree about the extent to which efficiency should be sacrificed to combat inequality.¹ Voters may be motivated by both their own self-interest and their views on what constitutes an equitable distribution, and it may be difficult to tease apart these two competing motivations. For example, in the United States, we typically associate the Democratic Party with the promotion of policies which reduce inequality, and the Republican Party with the promotion of efficiency. However, whether Democratic voters are more willing

¹In a classic series of writings, John Rawls and John Harsanyi argue that a “fair-minded” person must make distributive decisions that satisfy “the impartiality and impersonality requirements to the fullest possible degree” (Harsanyi 1978, p. 227) — in other words, the fair-minded should place equal weight on themselves and others. Harsanyi (1955) and Rawls (1971) nonetheless came to quite different conclusions about the equality-efficiency tradeoffs that fair-minded people should make in their distributional preferences. In fact, their familiar philosophical theories of distributive justice – utilitarianism and Rawlsianism – instill competing conceptions. Stated simply, Harsanyi argued that distributional preferences should maximize efficiency (increasing total payoffs), whereas Rawls argued that they should minimize inequity (reducing differences in payoffs).

to sacrifice efficiency — and even their own income — to reduce inequality is an open question; alternatively, Democrats may be those who expect to benefit from government redistribution, as the median voter theorem would suggest, or those who agree with other elements of the party’s platform. This highlights the importance of correctly distinguishing fair-mindedness from preferences over equality-efficiency tradeoffs and accurately measuring both in a large and diverse sample of American voters — in order to assess the extent to which distributional preferences explain political choices.

To this end, we conduct an incentivized experiment that is designed to measure fair-mindedness and equality-efficiency tradeoffs, using the American Life Panel (ALP), a longitudinal survey administered online by the RAND Corporation. The ALP sample consists of more than 5,000 individuals aged 18 and over recruited from several sources, including representative samples of the U.S. population. The ALP makes it possible to conduct sophisticated experiments via the internet, and to combine data from these experiments with detailed individual demographic and economic information. We invited a random sample of ALP respondents to participate in an incentivized online experiment involving real financial tradeoffs between oneself and another American; this allows us to examine the linkage between experimentally-elicited distributional preferences and the political decisions of Americans.

In our experiment, we study a modified two-person dictator game in which the set of monetary payoffs is given by the budget line $p_s\pi_s + p_o\pi_o = 1$, where π_s and π_o correspond to the payoffs of *self* (the subject) and an unknown *other* (an anonymous ALP respondent not sampled for the experiment), and $p = p_o/p_s$ is the relative price of redistribution.² This design allows us to decompose distributional preferences into

²The modified dictator game was first used by Andreoni and Miller (2002) and further developed by Fisman, Kariv and Markovits (2007), who introduced a graphical interface that makes it possible to present each subject with many choices in the course of a single experimental session. Using this graphical interface allows us to analyze behavior at the level of the individual subject, without the need to pool data or assume that subjects are

fair-mindedness and equality-efficiency tradeoffs: increasing the fraction of the budget spent on *other*, $p_o\pi_o$, as p increases indicates distributional preferences weighted towards equality (reducing differences in payoffs), whereas decreasing $p_o\pi_o$ when the relative price of redistribution increases indicates distributional preferences weighted towards efficiency (increasing total payoffs).

We begin our analysis of the experimental data by using revealed preference theory to determine whether observed choices are consistent with utility maximization. Because our subjects faced a wide range of intersecting budget lines, our data provide a stringent test of utility maximization. Although individual behaviors are complex and heterogeneous, we find that most subjects come close to satisfying the utility maximization model according to a number of standard measures. We therefore conclude that, at least in a controlled experimental setting where the tradeoffs are sufficiently transparent, most Americans are capable of making coherent and purposeful redistributive choices in the sense that these choices achieve a well-defined objective.

The consistency of individual decisions naturally leads us to ask what kind of distributional preferences are consistent with the observed choices. Our sample exhibits considerable heterogeneity in preferences, but relatively few subjects made choices that correspond to prototypical distributional preferences. Of our 1,002 subjects, 85 (8.5 percent) made choices consistent with Rawlsian distributional preferences, equalizing the payoffs to *self* and *other* regardless of the relative price of redistribution. Only two subjects displayed utilitarian preferences, maximizing the sum of the payoffs to *self* and *other* irrespective of the *ex post* inequality that would entail. Finally, only 81 subjects (8.1 percent) behaved selfishly, allocating themselves more than 95 percent of the total payoff, on average. These are, of course, special cases where the regularities in the data are very clear. To explain the distinct types of

homogenous.

individual behavior revealed by the full data set, we must impose further structure on the data.

To this end, we estimate individual-level utility functions of the constant elasticity of substitution (CES) form commonly employed in demand analysis. In the context of redistribution, the CES has the form

$$u_s(\pi_s, \pi_o) = [\alpha\pi_s^\rho + (1 - \alpha)\pi_o^\rho]^{1/\rho}$$

where α represents the degree of fair-mindedness (the relative weight on *self* versus *other*) and ρ characterizes equality-efficiency tradeoffs (the curvature of the altruistic indifference curves). Any $0 < \rho \leq 1$ indicates distributional preference weighted towards increasing total payoffs, whereas any $\rho < 0$ indicates distributional preference weighted towards reducing differences in payoffs. Our analysis generates individual-level estimates of $\hat{\alpha}_n$ and $\hat{\rho}_n$, allowing us to classify each subject's degree of fair-mindedness and equality-efficiency tradeoffs.

The estimation results for the CES specification reinforce the conclusion that distributional preferences vary widely across subjects. Table 1 provides a population-level summary of the parameter estimates. We classify subjects as either fair-minded, intermediate, or selfish, and as either equality-focused or efficiency-focused. 307 subjects (30.6 percent) are fair-minded, placing equal weight on the payoffs to *self* and *other*; while 161 subjects (16.0 percent) are selfish. Thus, fair-minded subjects outnumber selfish ones by about 2 to 1. 585 subjects (58.4 percent) are equality-focused, spending more on tokens for *other* when the relative price of redistribution is higher. In addition, we observe a greater degree of efficiency-focus among fair-minded subjects than among subjects who are more selfish.

Table 1 about here.

Exploiting the detailed demographic and economic data available on ALP sub-

jects, we then examine the correlates of the estimated CES parameters, $\hat{\alpha}_n$ and $\hat{\rho}_n$. Less educated subjects, as well as African Americans, are notably more fair-minded than the rest of the sample. Younger and lower income subjects and African Americans display greater efficiency focus, while women show greater equality focus. While observable attributes have predictive power in the data, we find that marked heterogeneity in distributional preferences remains within each demographic and economic group: observable attributes explain only about five percent of the variation in CES parameters.

Finally, and most importantly, after controlling for demographic characteristics and state of residence fixed effects, we find that our measure of efficiency focus, $\hat{\rho}_n$, is negatively related to the probability of having voted for Barack Obama in 2012, and also negatively related to the probability of reporting an affiliation with the Democratic Party. These results indicate that American voters are motivated by their distributional preferences governing equality-efficiency tradeoffs. By contrast, we do not find a significant relationship between our experimental measure of fair-mindedness, $\hat{\alpha}_n$, and either voting behavior or party affiliation; nor do we find that less fair-minded individuals from low (resp. high) income households are more likely to affiliate with the Democrats (resp. Republicans).³ These findings may be useful in explaining, in particular, the muted response to increased inequality in America. In the canonical median voter model of Meltzer and Richard (1981), an increasing skewness to the income distribution should increase the median voter's desire for equality-inducing redistribution. Yet if, as we find here, low-income voters have a stronger efficiency orientation, it may serve to counterbalance the increased demand for redistribution that the median voter theorem predicts would result from greater

³Because our measure of household income provides only a rough indicator of the likely beneficiaries of government redistribution, we do not view our results as evidence that self-interest plays no role in political decisions.

inequality.⁴

Overall, our findings contribute to the discussion around tax policy and other forms of government redistribution. In a standard model of taxation (Mirrlees 1971), moral hazard is the primary reason for incomplete redistribution, but standard estimates of labor-supply elasticity appear to predict much higher top income tax rates than are observed in modern developed economies (Diamond and Saez 2011). This has led scholars to propose a number of further explanations for the limited demand for inequality-reducing redistribution.⁵ In this paper, we provide a further, heretofore unexplored possibility: that Americans — in particular lower-income ones — have distributional preferences that emphasize efficiency over equality.

The fact that we find that our distributional preferences predict political decisions further strengthens the link between our findings and tax policy outcomes. As Saez and Stantcheva (2013) emphasize, optimal tax policy will depend on the distributional preferences of voters and taxpayers, and our work provides a first step in characterizing these preferences. Our design is particularly well-suited to this task, as subjects make tradeoffs between their own payoff and the payoff an individual drawn from the general population of the U.S. (another ALP respondent). This stands in contrast to many experiments, where subjects are generally matched with someone from their own community. Further, our experimentally generated measure of distributional preferences is not confounded by subjects' attitudes toward government in general, as is the case for survey-based measures of distributional preferences based on attitudes toward government redistribution (Saez and Stantcheva 2013).

The rest of the paper is organized as follows. The next section summarizes the

⁴This would involve augmenting voters' political preferences to include the income distribution itself, in such a way that efficiency orientation is negatively correlated with own income.

⁵See, for example, Benabou and Ok (2001) for the role of upward mobility; Lee and Roemer (2006) for the effects of “policy-bundling” redistribution with other, cross-cutting issues, and Norton and Ariely (2011) for the role of misinformation on the actual income distribution.

closely related literature. Section 3 describes the subject pool and the experimental design and procedures. Section 4 summarizes some important features of the individual-level data, and Section 5 describes the linkage between distributional preferences and political decisions. Section 6 contains some concluding remarks.

2 Related Literature

Experimental research has been very fruitful in documenting the existence of (non-selfish) distributional preferences and directing theoretical attention toward such preferences. We will not attempt to review the large and growing body of research on the topic. Key contributions include Loewenstein, Thompson and Bazerman (1989), Bolton (1991), Rabin (1993), Levine (1998), Fehr and Schmidt (1999), Bolton and Ockenfels (1998, 2000), Charness and Rabin (2002, 2005), and Andreoni and Miller (2002) among others. Camerer (2003) provides a comprehensive (if now somewhat dated) discussion of experimental and theoretical work in economics focusing on dictator, ultimatum, and trust games. The overarching lesson from hundreds of experiments is that people often sacrifice their own payoffs in order to increase the payoffs of (unknown) others, and they do so even in circumstances that do not engage reciprocity motivations or strategic considerations.

Fisman et al. (2007) extend the modified dictator game first proposed by Andreoni and Miller (2002), introducing an experimental technique (a graphical computer interface) that allows for the collection of richer individual-level data from dictator game experiments than had previously been possible. This is particularly important given that, as Andreoni and Miller (2002) emphasize, individual preferences are heterogeneous, so behavior must be examined at the individual level for distributional preferences to be properly understood. Fisman, Jakiela, Kariv and Markovits (2015b) demonstrate the predictive validity of the preference parameters elicited using Fisman et al.'s (2007) graphical dictator game interface by showing that our experimental

measure of equality-efficiency tradeoffs predicts the subsequent career choices of Yale Law School students — more efficiency-focused students are more likely choose careers in corporate law, while more equality-focused students are more likely to work in the non-profit sector.⁶

This paper is most closely related to other distributional preference experiments that have used subjects drawn from broad cross-sections of the adult population (as opposed to university students). Bellemare, Kröger and van Soest (2008) study distributional preferences in a large and heterogeneous sample of Dutch adults. In their experiment, survey respondents from the CentERpanel participate in ultimatum games. Like the ALP, the CentERpanel implements sophisticated experiments and collects extensive demographic and economic information from its members. Data characterizing subjects' decisions within the experiment, their beliefs about the likelihood that specific ultimatum game offers would be accepted, and their individual characteristics are used to estimate a structural model of inequality aversion (Fehr and Schmidt 1999) in the Dutch population. By comparison, we restrict attention to dictator games, which allows us to focus on behavior motivated by purely distributional preferences and thus ignore the complications of strategic behavior and reciprocity motivations inherent in response games.⁷

Our overall findings resonate with those of Bellemare et al. (2008) — who also find considerable heterogeneity in preferences, much of which is not correlated with

⁶In related work, Fisman, Jakiela and Kariv (2015a) use the same experimental methodology to estimate the impact of the Great Recession on distributional preferences.

⁷While Bellemare et al. (2008) also conduct dictator games, they only use decisions in those games to assess the predictive power of the structural parameter estimates derived from ultimatum game decisions. Bellemare, Kröger and van Soest (2011) combine the data from Bellemare et al. (2008) with data on responders in a (random) ultimatum game in order to separate distributive concerns from the intentions subjects attribute to the actions of others. In another experimental paper on distributional preferences with the CentERpanel, Bellemare and Kröger (2007) use an investment game that builds on Berg, Dickhaut and McCabe (1995) to study the correlations between distributional preferences and demographic and socioeconomic characteristics.

observable characteristics — but there are several key differences between the two studies, beyond the fact that we draw our samples from different societies. First, our experiment allows us to explicitly test whether individual choices can be rationalized by a utility function defined over payoffs to *self* and *other*. To our knowledge, no such tests of the rationality of individual distributional preferences have been conducted in the general population. Second, our experimental design allows us to separately identify fair-mindedness and equality-efficiency tradeoffs, and to estimate individual utility functions at the subject level. Their study makes more restrictive assumptions about the functional form of the utility function and the distribution of unobservable heterogeneity within the population. Finally, Bellemare et al. (2008) explore the relationship between beliefs (specifically, optimism about others' fair-mindedness) and distributional preferences, while we focus on the relationship between equality-efficiency tradeoffs measured in the laboratory and political decisions in the real world. Our findings may thus be used to enrich models of voting and/or political competition, and additionally add to our understanding of policy formation in the U.S.

Like Bellemare et al. (2008, 2011), our work also contributes to the rapidly expanding literature characterizing the distributional preferences of the general (non-student) population. Much of this work focuses on cross-country differences in distributional preferences; seminal contributions include Roth, Prasnikar, Okuno-Fujiwara and Zamir (1991), Henrich, McElreath, Barr, Ensminger, Barrett, Bolyanatz, Cardenas, Gurven, Gwako, Henrich, Lesorogol, Marlowe, Tracer and Ziker (2006), and Henrich, Ensminger, McElreath, Barr, Barrett, Bolyanatz, Cardenas, Gurven, Gwako, Henrich, Lesorogol, Marlowe, Tracer and Ziker (2010). Our work is most closely related to papers such as Hermann, Thöni and Gächter (2008) that explore the connections between the distributional preferences of a population and political economic outcomes within that country.

3 Experimental Design

3.1 Subject Pool

We embed an incentivized experiment in the American Life Panel (ALP), an internet survey administered by the RAND Corporation to more than 5,000 adult Americans. ALP respondents have been recruited in several different ways, including from representative samples of the U.S. population.⁸ To recruit subjects for our experiment, ALP administrators sent email invitations to a random sample of ALP respondents. 1,172 ALP respondents received the email and logged in to the experiment.⁹ Of those, 1,043 (89.9 percent) progressed to the incentivized decision problems and 1,002 respondents (85.5 percent) completed the entire experiment; these subjects constitute our subject pool.

Table 2 compares the ALP sample to the American Community Survey (ACS) conducted by the U.S. Census and representative of the U.S. population in 2012. We present the data for participants (those who completed the experiment); participants plus dropouts (those who logged in but then quit the experiment); participants, dropouts, and non-participants (those who were invited to participate in the experiment but never logged in); and the entire ALP sample. Like the U.S. population, the ALP sample includes an enormous amount of demographic, socioeconomic, and

⁸The initial participants were selected from the Monthly Survey Sample of the University of Michigan’s Survey Research Center. Additional respondents have been added through random digit dialling, targeted recruitment of a vulnerable population sample of low-income individuals, and snowball sampling of existing panel members. See the ALP website (<https://mmicdata.rand.org/alp/>) for information on panel composition, demographics, attrition and response rates, sampling weights, and a comparison with other data sources.

⁹Those ALP respondents for whom complete demographic information was unavailable were not eligible to participate. ALP administrators sent email invitations to a random sample of 1,700 respondents (out of approximately 4,000) for whom a valid email address and complete demographic information was available. We are unable to distinguish subjects who read the invitation email and chose not to participate from those who never received the invitation (for example, because they do not regularly access the email account registered with the ALP).

geographic diversity; the subsample of 1,002 ALP respondents that constitute our subject pool is remarkably consistent with the entire ALP sample.¹⁰

Table 2 about here.

Subjects in our experiment are from 47 U.S. states, and range in age from 19 to 91. 58 percent are female. 9 percent of our subjects did not finish high school, while 31 percent hold college degrees. 56 percent of subjects are currently employed; the remainder include retirees (17 percent), the unemployed (11 percent), the disabled (8 percent), homemakers (6 percent), and others who are on medical leave or otherwise temporarily absent from the workforce. 68 percent identify themselves as non-Hispanic whites, 18 percent as Hispanic or Latino, and 11 percent as African American. 18 percent live in the Northeast (census region I), 20 percent in the Midwest (census region II), 35 percent in the South (census region III), and 267 percent in the West (census region IV). Our subject pool therefore contains under-represented groups in terms of age, educational attainment, household income, occupational status, and place of residence.

3.2 Experimental Procedures

To provide a positive account of individual distributional preferences, one needs a choice environment that is rich enough to allow a general characterization of patterns of behavior; Fisman et al. (2007) developed a computer interface for exactly this purpose. The interface presents a standard consumer decision problem as a graphical representation of a budget line and allows the subject to make choices through a simple point-and-click design.¹¹

¹⁰In the Online Appendix, we examine the individual characteristics associated with completing the experiment. Those who complete the experiment look similar to the overall ALP subject pool and the subset of ALP respondents invited to participate in the experiment.

¹¹The experimental method is applicable to many types of individual choice problems. See Choi, Fisman, Gale and Kariv (2007) and Ahn, Choi, Gale and Kariv (2014), for

In this paper, we study a modified dictator game in which a subject divides an endowment between *self* and an anonymous *other*, an individual chosen at random from among the ALP respondents not sampled for the experiment. The subject is free to allocate a unit endowment in any way she wishes subject to the budget constraint, $p_s\pi_s + p_o\pi_o = 1$, where π_s and π_o denote the payoffs to *self* and *other*, respectively, and $p = p_o/p_s$ is the relative price of redistribution. This decision problem is presented graphically on a computer screen, and the subject must choose a payoff allocation, (π_s, π_o) , from a budget line representing feasible payoffs to *self* and *other*.¹² Confronting subjects with a rich menu of such budget lines allows us to identify both the tradeoff between both *self* and *other* (fair-mindedness) and the tradeoff between equality and efficiency — because responses to price changes allow us to separately identify these tradeoffs. A subject who increases the fraction of the budget spent on *other* as the relative price of redistribution increases has preferences weighted towards equality (i.e. minimizing differences in payoffs), while a subject who decreases the fraction of the budget spent on *other* as the relative price of redistribution increases has preferences weighted towards efficiency (maximizing the aggregate payoff).¹³

The experiment consisted of 50 independent decision problems. For each decision problem, the computer program selected a budget line at random from the set of lines that intersect at least one of the axes at 50 or more experimental currency tokens, but with no intercept exceeding 100 tokens. Subjects made their choices by using the settings involving, respectively, risk and ambiguity. Choi, Kariv, Müller and Silverman (2014) investigate the correlation between individual behavior under risk and demographic and economic characteristics within the CentERpanel, a representative sample of more than 2,000 Dutch households; that project demonstrated the feasibility of using the graphical experimental interface in web-based surveys.

¹²Full experimental instructions are included in the Online Appendix.

¹³In a standard dictator experiment (cf. Forsythe, Horowitz, Savin and Sefton 1994), $\pi_s + \pi_o = 1$: the set of feasible payoff pairs is the line with a slope of -1 , so the problem is simply dividing a fixed total income between *self* and *other*, and there is no inherent tradeoff between equality and efficiency.

computer mouse or keyboard arrows to move the pointer to the desired allocation, (π_s, π_o) , and then clicked the mouse or hit the enter key to confirm their choice.

At the end of the experiment, payoffs were determined as follows. The experimental program first randomly selected one of the 50 decision problems to carry out for real payoffs. Each decision problem had an equal probability of being chosen. Each subject then received the tokens that she allocated to *self* in that round, π_s , while the randomly-chosen ALP respondent with whom she was matched received the tokens that she allocated to *other*, π_o .¹⁴ Payoffs were calculated in terms of tokens and then translated into dollars at the end of the experiment. Each token was worth 50 cents. Subjects received their payments from the ALP reimbursement system via direct deposit into a bank account.

4 Decomposing Distributional Preferences

One aspect of the rich data generated by the experiment is that they allow us to analyze behavior at the level of individual subjects, testing whether choices are consistent with individual utility maximization and if so identifying the structural properties of the underlying utility function, without the need to pool data or assume that subjects are homogenous. If budget sets are linear (as in our experiment), classi-

¹⁴To describe preferences with precision at the individual level, it is necessary to generate many observations per subject over a wide range of budget sets. Our subjects made decisions over 50 budget sets, with one decision round selected at random from each subject to carry out for payoffs. This random selection approach is a standard practice, although it is the subject of ongoing controversy in the literature. If we paid for all rounds, subjects could easily hedge against inequality. The random payoff method prevents such hedging and reveals underlying distributional preferences only under stringent independence conditions. However, hedging relies heavily on the fact that the individual knows the parameters of future budget set. In our experiment, subjects faced a large menu of highly heterogeneous budget sets, and were only informed about the price's random generating process, making it difficult to hedge. Finally, given the novelty of our experimental design, we wished to keep as many aspects of the experiment consistent with prior studies as was possible. The random selection approach is the method used by Andreoni and Miller (2002), among many others.

cal revealed preference theory (Afriat (1967); Varian (1982, 1983)) provides a direct test: choices in a finite collection of budget sets are consistent with maximizing a well-behaved utility function if and only if they satisfy the Generalized Axiom of Revealed Preference (GARP). To account for the possibility of errors, we assess how nearly individual choice behavior complies with GARP by using Afriat’s (1972) Critical Cost Efficiency Index (CCEI). We find that most subjects exhibit GARP violations that are minor enough to ignore for the purposes of recovering distributional preferences or constructing appropriate utility functions. To economize on space, the revealed preference analysis is provided in an Online Appendix, where we also analyze the experimental data using a reduced form approach that imposes no functional form assumptions on distributional preferences. The estimations presented in the paper based on the CES utility specification convey the same message as the estimates from the reduced form framework.

4.1 The CES Utility Specification

Our subjects’ CCEI scores are sufficiently close to one to justify treating the data as utility-generated, and Afriat’s theorem tells us that the underlying utility function, $u_s(\pi_s, \pi_o)$, that rationalizes the data can be chosen to be increasing, continuous and concave. In the case of two goods, consistency and budget balancedness imply that demand functions must be homogeneous of degree zero. If we also assume separability and homotheticity, then the underlying utility function, $u_s(\pi_s, \pi_o)$, must be a member of the constant elasticity of substitution (CES) family commonly employed in demand analysis:

$$u_s(\pi_s, \pi_o) = [\alpha(\pi_s)^\rho + (1 - \alpha)(\pi_o)^\rho]^{1/\rho} \quad (1)$$

where $0 \leq \alpha \leq 1$ and $\rho \leq 1$.¹⁵ The CES specification is very flexible, spanning a range of well-behaved utility functions by means of the parameters α and ρ . The parameter α represents the weight on payoffs to *self* versus *other* (fair-mindedness), while ρ parameterizes the curvature of the indifference curves (equality-efficiency tradeoffs).

When $\alpha = 1/2$, a subject is fair-minded in the sense that *self* and *other* are treated symmetrically. Among fair-minded subjects, the family of CES utility functions spans the spectrum from Rawlsianism to utilitarianism as ρ ranges from $-\infty$ to 1. In particular, as ρ approaches $-\infty$, $u(\pi_s, \pi_o)$ approaches $\min\{\pi_s, \pi_o\}$, the maximin utility function of a Rawlsian; as ρ approaches 1, $u(\pi_s, \pi_o)$ approaches that of a utilitarian, $\pi_s + \pi_o$. Hence, both the Rawlsian and the utilitarian utility functions, as well as a whole class of intermediate fair-minded utility functions, are admitted by the CES specification.¹⁶

¹⁵The proper development of revealed preference methods to test whether data are consistent with a utility function with some special structure, particularly homotheticity and separability, is beyond the scope of this paper. Varian (1982, 1983) provides combinatorial conditions that are necessary and sufficient for extending Afriat's (1967) Theorem to testing for special structure of utility, but these conditions are not simple adjustments of the usual tests, which are all computationally intensive for large datasets like our own.

¹⁶In Figure 1 we depict a budget line where $p_s > p_o$ (the relative price of redistribution is less than 1) and highlight the allocations consistent with prototypical fair-minded distributional preferences. The point *A*, which lies on the diagonal, corresponds to the equal allocation $\pi_s = \pi_o$. This allocation is consistent with Rawlsian or maximin distributional preferences, which are characterized by right-angle indifference curves (and the utility function $u_s(\pi_s, \pi_o) = \min\{\pi_s, \pi_o\}$). Point *B* represents an allocation in which $\pi_s = 0$ and $\pi_o = 1/p_o$, consistent with the linear indifference curves characterizing utilitarian preferences (with the utility function $u_s(\pi_s, \pi_o) = \pi_s + \pi_o$). The Rawlsian and utilitarian preferences represent the two ends of the spectrum of equality-efficiency tradeoffs. The centroid of a budget line, *C*, represents an allocation with equal budget shares spent on *self* and *other* such that $p_s\pi_s = p_o\pi_o$. This allocation is consistent with Cobb-Douglas preferences (characterized by the utility function $u_s(\pi_s, \pi_o) = \pi_s\pi_o$). In this case the equality-efficiency tradeoffs are intermediate between Rawlsian and utilitarian preferences. More generally, the concavity of $u_s(\pi_s, \pi_o)$ measures aversion to inequality. Finally, note that because the distributional preferences depicted in Figure 1 are fair-minded, each indifference curve is symmetric with respect to the diagonal. Increasing the weight on *self* relative to *other* shifts indifference curves upwards.

Figure 1 about here.

More generally, different values of ρ give different degrees to which equality is valued over efficiency. Any $0 < \rho \leq 1$ indicates distributional preference weighted towards efficiency (increasing total payoffs) because the expenditure on the tokens given to *other*, $p_o\pi_o$, decreases when the relative price of giving $p = p_o/p_s$ increases, whereas any $\rho < 0$ indicates distributional preference weighted toward equality (reducing differences in payoffs) because $p_o\pi_o$ increases when p increases. As ρ approaches 0, $u(\pi_s, \pi_o)$ approaches the Cobb-Douglas utility function, $\pi_s^\alpha \pi_o^{1-\alpha}$, so the expenditures on tokens to *self* and *other* are constant for any price p — a share α is spent on tokens for *self* and a share $1 - \alpha$ is spent on tokens for *other*.

The CES expenditure function is given by

$$p_s \pi_s = \frac{g}{(p_s/p_o)^r + g}$$

where $r = \rho/(\rho - 1)$ and $g = [\alpha/(1 - \alpha)]^{1/(1-\rho)}$. This generates the following individual-level econometric specification for each subject n :

$$p_{s,n}^i \pi_{s,n}^i = \frac{g_n}{(p_{s,n}^i/p_{o,n}^i)^{r_n} + g_n} + \epsilon_n^i$$

where $i = 1, \dots, 50$ indexes the decision round and ϵ_n^i is assumed to be distributed normally with mean zero and variance σ_n^2 . We normalize prices at each observation and estimate demand in terms of expenditure shares, which are bounded between zero and one, with an *i.i.d.* error term.¹⁷ We generate individual-level estimates \hat{g}_n and \hat{r}_n using non-linear Tobit maximum likelihood, and use these estimates to infer

¹⁷For perfectly consistent subjects, there exists a (well-behaved) utility function that choices maximize (as implied by Afriat's Theorem) so the error term in our individual-level regression analysis can only stem from misspecifications of the functional form. For less than perfectly consistent subjects, the error term also captures the fact these subjects compute incorrectly, execute intended choices incorrectly, or err in other ways. Disentangling these sources of noise is beyond the scope of this paper.

the values of the underlying CES parameters $\hat{\alpha}_n$ and $\hat{\rho}_n$.

4.2 Fair-mindedness

We classify a subject as fair-minded if $\hat{\alpha}_n$ is between 0.45 and 0.55. By this criterion, 307 subjects (30.6 percent) are fair-minded. In contrast, only 161 subjects (16.0 percent) are selfish ($\hat{\alpha}_n > 0.95$). Thus, fair-minded subjects outnumber selfish ones by approximately 2 to 1.¹⁸

Figure 2 explores the extent to which heterogeneity in estimated $\hat{\alpha}_n$ parameters is explained by observable characteristics. We partition the subject pool into mutually exclusive categories to examine variation by age, gender, education, and so forth. The means for all categories are clustered between 0.6 and 0.71. The averages suggest that fair-mindedness generally decreases with age, education, and household income. In particular, subjects with less than a high school diploma are particularly fair-minded; the unemployed and, to a lesser extent, the disabled appear more fair-minded than employed subjects, retirees, and homemakers. Consistent with other studies (cf. Croson and Gneezy 2009), we find that women are more fair-minded than men, though the effect is quite small.¹⁹ We also find that non-Hispanic whites are significantly less fair-minded than both African American and Hispanic subjects.

Figure 2 about here.

In Table 3, we explore the associations between fair-mindedness and individual characteristics in a regression framework.²⁰ We report the results of OLS regressions of $\hat{\alpha}_n$ on our full set of individual characteristics; in Column 2, we also include state

¹⁸We obtain similar results using other thresholds to identify subjects' types, or if we use statistical tests to classify types.

¹⁹The mean $\hat{\alpha}_n$ is 0.67 among women and 0.69 among men (p-value 0.06). However, women are 4.5 percentage points more likely to be fair-minded (p-value 0.124).

²⁰We report OLS regression results, but findings are unchanged if we adopt a Tobit specification to account for censoring of $\hat{\alpha}_n$ at 1. The results are nearly identical because very few subjects have very high estimated $\hat{\alpha}_n$ parameters.

of residence fixed effects. The indicators for being African American and having less than a high school education are both negative and significant in both specifications, indicating that these groups are, on the whole, more fair-minded than other Americans. We again find that much of the observed heterogeneity in fair-mindedness occurs within rather than across groups: demographic and socioeconomic characteristics explain only 4.1 percent of the variation in $\hat{\alpha}_n$.²¹

Table 3 about here.

4.3 Equality-Efficiency Tradeoffs

The mean $\hat{\rho}_n$ observed in our sample is -2.64 , and the median is -0.184 . 585 subjects (58.4 percent) have estimated $\hat{\rho}_n$ parameters below 0, indicating a greater focus on equality (minimizing differences in payoffs) than on efficiency (maximizing total payoffs). In Figures 3 and 4, we disaggregate the estimated $\hat{\rho}_n$ parameters by demographic and socioeconomic categories (Figure 3 presents means and Figure 4 presents medians). Three main results stand out. First, the youngest subjects are substantially more efficiency-focused than all of the three older quartiles. The median $\hat{\rho}_n$ among subjects in the youngest quartile is 0.025, while the median in older quartiles is -0.276 . Second, non-Hispanic whites are substantially less efficiency-focused than minorities. The median $\hat{\rho}_n$ is -0.321 among non-Hispanic whites in our sample, while the medians for Hispanic and African American subjects are -0.037 and 0.092, respectively. Finally, subjects from low income households are more efficiency-focused than wealthier individuals. This last finding may help to explain the fact that the increase in income inequality observed in the United States in recent decades has not led to increased political support for redistributive policies.

Figures 3 and 4 about here.

²¹Adding state fixed effects raises the amount of variation that is explained by observables to 8.9 percent.

In Table 4, we explore the associations between equality-efficiency tradeoffs and individual characteristics in a regression framework. In Columns 1 and 3, we replicate the OLS specifications from Table 3 with $\hat{\rho}_n$ as the dependent variable. Since estimates of $\hat{\rho}_n$ are quite noisy for relatively selfish subjects, we also report (in the even-numbered columns) specifications which omit the 45 subjects who allocate themselves an average of more than 99 percent of tokens. Given the skewed distribution of the estimated $\hat{\rho}_n$ parameters, we also report several alternative specifications: we report median regressions in Columns 5 and 6, regressions in which the outcome variable is the decile of the estimated $\hat{\rho}_n$ distribution in Columns 7 and 8, and regressions in which the dependent variable is an indicator for having $\hat{\rho}_n \geq 0$ (which we term ρ_{high}) in Columns 9 and 10.

Table 4 about here.

Several robust associations, already hinted at by the patterns in Figures 3 and 4, stand out. First, the youngest quartile of subjects are significantly more efficiency-focused than older individuals (in all specifications). The coefficient in Column 10 of Table 4, for example, suggests that subjects in the youngest quartile are 8.1 percentage points more likely to be focused on efficiency in the sense of having $\hat{\rho}_n$ of at least 0. Second, subjects with household incomes in the lowest income quartile are also significantly more focused on efficiency than the rest of the sample. Third, women are significantly more focused on equality than men (a finding consistent with the evidence reported in Andreoni and Vesterlund (2001)). Finally, though the association is not significant in all specifications, our results suggest that African American subjects are more efficiency-focused than non-Hispanic whites (the omitted category). Although this coefficient is significant in only 6 of the 10 specifications, the point estimates are extremely large, suggesting, for example, that African Americans are 17.6 percentage points more likely to have a $\hat{\rho}_n$ of at least 0. As with fair-mindedness, we find that much of the observed variation in equality-efficiency preferences occurs

within rather than between groups. Individual demographic and socioeconomic characteristics explain 4.36 percent of the variation in $\hat{\rho}_n$ in our sample. Thus, though some groups appear more efficiency-focused than others, these between-group differences are modest relative to the tremendous variation in efficiency orientation within the demographic and socioeconomic categories in our sample.

5 Distributional Preferences and Political Behavior

In our final piece of analysis, we test whether distributional preferences, as measured in our experiment, predict support for political candidates who favor redistribution. Whether efficiency-focused distributional preferences are associated with political support for government redistribution is an open question. A vast literature on the partisan preferences of Americans assumes that Democrats have stronger preferences for inequality-reducing government policy than Republicans, a view that is validated based on survey responses to the General Social Survey (Hayes 2011). However, as Kuziemko, Norton, Saez and Stantcheva (2013) point out, this does not necessarily imply that Democrats are more averse to inequality; they may instead look more favorably on government intervention in general, and on redistributive policies in particular. Indeed, when Kuziemko et al. (2013) remove government involvement from questions regarding inequality, they find that much of the partisan difference in distributional preferences disappears.

Our experiment provides an objective measure of the extent to which individuals actually choose to sacrifice efficiency to reduce inequality (when matched with another randomly-chosen American), an approach which contrasts markedly with research designs based on non-incentivized survey questions. Further, our measure of equality orientation is removed from any association between redistribution and government intervention. To the best of our knowledge, there is no empirical evidence on whether the equality-efficiency tradeoffs elicited through such incentivized lab experiments

predict voting behavior.²²

We explore the link between equality-efficiency tradeoffs and political behavior by looking at voting decisions in the 2012 presidential election.²³ Our main dependent variable is an indicator for voting for Democrat Barack Obama, a relatively pro-redistribution candidate, rather than Republican Mitt Romney.²⁴ We focus on the 766 subjects who participated in ALP modules exploring participants' choices in the 2012 election and who report voting for either Barack Obama or Mitt Romney.²⁵ We include a range of demographic controls to account for the fact that, for example, African Americans overwhelmingly voted for Obama for reasons that are plausibly distinct from their distributional preferences.²⁶ We employ a linear probability model with an indicator variable for having voted for Obama as the outcome.²⁷ We report results for all three measures of equality-efficiency tradeoffs used in the preceding

²²Durante, Putterman and van der Weele (forthcoming) find that more politically liberal university students support higher within-experiment tax rates. They present a cleverly designed experiment that distinguishes between three motives for supporting income redistribution — own-income maximization (of those in lower income brackets), risk aversion, and distributional preferences. Subjects in these experiments are undergraduate students at Brown University. Their conclusion is that own-payoff maximization is the dominant motive for redistribution in the experiment, but distributional preferences also play a key role in subjects' decisions.

²³Data on voting behavior in the 2008 election is not available for most of our subjects, in part because the ALP sample is regularly refreshed with new respondents, and because most studies recruit only a small fraction of ALP respondents (so the overlap between our randomly-chosen subjects and those who participated in other studies is limited).

²⁴As one indication of their views on redistribution, in September 2012, media outlets reported the discovery of a recording of Barack Obama (from 1998) stating that he “actually believe[d] in redistribution.” In response to the media coverage of the recording, Mitt Romney indicated that he “disagree[d].”

²⁵Unfortunately, no information is available on the voting behavior of the 48 subjects who participated in the relevant ALP survey module but did not report casting a ballot for a major party candidate, so we cannot classify the candidates that they supported as being either for or against redistribution.

²⁶Interestingly, without controls, the relationship between measured distributional preferences and voting is insignificant in all regressions, reflecting the fact that groups such as African Americans and low income individuals tend to support Democratic candidates, but are also more efficiency-focused in our experiments.

²⁷Probit results are nearly identical.

section: the estimated $\hat{\rho}_n$ parameter; $\hat{\rho}_n$ deciles; and ρ_{high} , an indicator for being efficiency-focused in the sense of having an estimated $\hat{\rho}_n$ of at least 0.

In the first three columns of Table 5, we present specifications which include only demographic controls, showing results for each of the three transformations of $\hat{\rho}_n$. The most straightforward coefficient to interpret is that on ρ_{high} in Column 3, which indicates that efficiency-focused subjects (with $\hat{\rho}_n \geq 0$) are 4.5 percentage points less likely to have voted for Obama than Romney. While the coefficients on $\hat{\rho}_n$ and its transformations are negative across the three columns, none is significant. The relationship is potentially confounded by large differences across regions in both equality-efficiency tradeoffs and voting patterns, however. For example, there is a strong equality orientation in Southern states, which also tend to vote Republican. In Columns 4 through 6 we repeat our analyses including state fixed-effects to absorb differences across geography. The coefficients on $\hat{\rho}_n$ and its variants are now significant at either the 5 or 10 percent level (p-values range from 0.02 to 0.07). The most easily interpreted coefficient, on ρ_{high} in Column 8, is -0.068. To provide a benchmark for the magnitude of this effect, we include (in the Online Appendix) the full set of regression coefficients from specifications with and without the inclusion of ρ_{high} as a covariate. We observe, for example, that the impact of ρ_{high} is greater than the effect of gender (0.054), and only marginally smaller than the impact of moving from low to medium (0.098) or medium to high (0.092) income. It is also of note that the coefficient on FEMALE declines somewhat with the inclusion of ρ_{high} , indicating that some amount of the gender voting gap can be directly accounted for by distributional preferences.

Table 5 about here.

In Panel B of Table 5, we omit nearly selfish subjects who allocate an average of more than 99 percent of the tokens to *self* because estimates of $\hat{\rho}_n$ are quite noisy for these individuals. All of our point estimates are marginally higher and the standard

errors unchanged, leading to marginally higher levels of statistical significance across all specifications.

We further explore the relationship between equality-efficiency tradeoffs and political behavior by replicating our specifications using an indicator for alignment with the Democratic Party as an outcome variable. These specifications include 528 subjects who participated in ALP modules on politics and identified themselves as either Republicans or Democrats.²⁸ We report our results in Table 6. All estimated coefficients on $\hat{\rho}_n$ and its transformations are negative, suggesting that more efficiency-focused subjects are less likely to be Democrats. Both the decile of $\hat{\rho}_n$ and, ρ_{high} the indicator for having $\hat{\rho}_n > 0$, are significant at least at the 90 percent level, with and without state fixed effects (p-values range from 0.051 to 0.003). After controlling for individual characteristics and geographic fixed effects, the estimated coefficient on $\hat{\rho}_{high}$ suggests that efficiency-focused subjects are 10.4 percentage points less likely to be Democrats. We again provide the full set of regression coefficients in the Online Appendix, both with and without ρ_{high} as a covariate. For the dependent variable of DEMOCRAT, the impact of $\hat{\rho}_n$ is greater than any covariate, apart from those related to race or religious affiliation.

Table 6 about here.

Overall, our results strongly suggest that the political decisions of Americans are motivated by their equality-efficiency preferences, and not just their own self-interest or their views of government. However, this pattern only emerges after one accounts for the fact that poorer Americans and minorities are, overall, substantially more focused on efficiency than the rest of the population.

²⁸Results are similar when we include the 217 additional subjects who participated in the politics module and identified themselves as Independents. 55 subjects participated in the module but indicated their party affiliation as “other,” so their parties cannot be classified as more or less equality-focused than the Democrats.

Finally, we also explore the relationship between fair-mindedness and political preferences. Paralleling our analysis of the link between equality-efficiency tradeoffs and political preferences, in the regression results reported in Table 7 the independent variable is an indicator for voting for Barack Obama in the 2012 while in Table 8 the independent variable is affiliation with the Democratic Party. Across all specifications, we find no significant relationship between our experimental measure of fair-mindedness, $\hat{\alpha}_n$, and either voting behavior or party affiliation across all specifications. This insignificant effect could be masking the opposing effects of self-interest on voting behavior for different sub-populations: a self-interested low-income individual should favor Democrats, while the opposite should be the case for a self-interested high-income individual.²⁹ Interestingly, we do not find support for this view in the data: the correlation between fair-mindedness and political preferences does not differ significantly by subject income. This suggests that Americans may not vote for redistributive policies purely out of (monetary) self-interest.

Table 7 about here.

6 Conclusion

In this paper, we take a first step in characterizing, via experiments administered through the ALP, the distributional preferences of the general population of the United States. While we observe a great deal of heterogeneity in the selfishness of subjects, we document a much higher rate of fair-mindedness than prior studies that involved primarily university students. There is also considerable heterogeneity in subjects' equality-efficiency tradeoffs. Some of the heterogeneity in subjects' distributional preferences can be explained by observable attributes, at times in un-

²⁹An extensive literature explores the extent to which voters support policies that are in their own perceived short-run and long-run economic interests. See, Alesina and La Ferrara (2005) and the references cited therein.

expected ways. Wealthier subjects, for example, are relatively less fair-minded; while low income subjects and African Americans are more focused on efficiency. But overall the data indicate a high degree of heterogeneity within each demographic or economic category. Most notably, we observe a strong relationship between our experimental measure of efficiency-equality tradeoffs and political decisions, thereby providing a link from underlying distributional preferences to voter preferences over policy outcomes.

These results are important for the formulation of a range of social and redistributive policies, and for understanding support for such policies. Distributional preferences are important inputs into any measure of social welfare — for example, optimal taxation hinges on an understanding of the distributional preferences of the population. Recent work in public finance (cf. Saez and Stantcheva 2013) also highlights the potential role of distributional preferences in explaining support for observed tax policies, which are not considered optimal from a theoretical perspective given standard assumptions about individual utilities.

Thus, our findings may be useful in providing a positive explanation of public opinion on policy issues related to redistribution. Most standard models of self-interested political preferences predict that the increase in income inequality observed in the United States over the last few decades should have led to greater support for government redistribution. However, no such shift has been observed in survey data (Kuziemko et al. 2013). Our findings partially explain this: voters are motivated by their distributional preferences, so they may not vote for redistributive policies which would make them better off individually.³⁰

³⁰Redistributive decisions depend on recipient attributes, and in particular ought to be a function of recipient income. In other work-in-progress, we study how distributional preferences vary based on the income of *other*, and with the degree of inequality between *self* and *other*. Introducing information about actual incomes into our experimental setup is perhaps the most important step toward understanding the distributional preferences most relevant to policy preferences and voting behavior, simply because views about how much income ought to be redistributed depend crucially on the initial incomes of the potential

Moreover, our results show that lower income Americans are more focused on efficiency than other groups; while, in related work, Kuziemko (in progress) and Fisman et al. (2015a) present suggestive evidence indicating that there may, in fact, be a causal relationship between negative income shocks and efficiency focus. Taken together, these results may help to explain why the increase in inequality observed in the United States has not led to any shift in party platforms toward greater redistribution in recent years.

recipients of redistribution.

References

- Afriat, Sydney N.**, “The Construction of Utility Functions from Expenditure Data,” *International Economic Review*, 1967, 8 (1), 67–77.
- , “Efficiency Estimates of Production Functions,” *International Economic Review*, 1972, 8, 568–598.
- Ahn, David, Syngjoo Choi, Douglas Gale, and Shachar Kariv**, “Estimating Ambiguity Aversion in a Portfolio Choice Experiment,” *Quantitative Economics*, 2014, 5 (2), 195–223.
- Alesina, Alberto and Eliana La Ferrara**, “Preferences for Redistribution in the Land of Opportunities,” *Journal of Public Economics*, 2005, 89 (5), 897–931.
- Andreoni, James and John Miller**, “Giving According to GARP: An Experimental Test of the Consistency of Preferences for Altruism,” *Econometrica*, 2002, 70 (2), 737–753.
- and **Lise Vesterlund**, “Which Is the Fairer Sex? Gender Differences in Altruism,” *Quarterly Journal of Economics*, 2001, 116 (1), 293–312.
- Bellemare, Charles and Sabine Kröger**, “On Representative Social Capital,” *European Economic Review*, 2007, 51 (1), 183–202.
- , —, and **Arthur van Soest**, “Measuring Inequity Aversion in a Heterogeneous Population Using Experimental Decisions and Subjective Probabilities,” *Econometrica*, 2008, 76 (4), 815–839.
- , —, and —, “Preferences, Intentions, and Expectation Violations: A Large-Scale Experiment with a Representative Subject Pool,” *Journal of Economic Behavior and Organization*, 2011, 78 (3), 349–365.

- Benabou, Roland and Efe A Ok**, “Social Mobility and the Demand for Redistribution: The Poupou Hypothesis,” *Quarterly Journal of Economics*, 2001, pp. 447–487.
- Berg, Joyce E., John Dickhaut, and Kevin McCabe**, “Trust, reciprocity, and social history,” *Games and Economic Behavior*, 1995, 10, 122–142.
- Bolton, Gary E**, “A Comparative Model of Bargaining: Theory and Evidence,” *The American Economic Review*, 1991, pp. 1096–1136.
- and **Axel Ockenfels**, “Strategy and Equity: an ERC-Analysis of the Güth–van Damme Game,” *Journal of Mathematical Psychology*, 1998, 42 (2), 215–226.
- and — , “ERC: A Theory of Equity, Reciprocity, and Competition,” *American Economic Review*, 2000, pp. 166–193.
- Camerer, Colin**, *Behavioral Game Theory*, Princeton, NJ: Princeton University Press, 2003.
- Charness, Gary and Matthew Rabin**, “Understanding Social Preferences with Simple Tests,” *Quarterly Journal of Economics*, 2002, 117 (3), 817–869.
- and — , “Expressed Preferences and Behavior in Experimental Games,” *Games and Economic Behavior*, 2005, 53 (2), 151–169.
- Choi, Syngjoo, Raymond Fisman, Douglas Gale, and Shachar Kariv**, “Consistency and Heterogeneity of Individual Behavior under Uncertainty,” *American Economic Review*, 2007, 97 (5), 1921–1938.
- , **Shachar Kariv, Wieland Müller, and Dan Silverman**, “Who Is (More) Rational?,” *American Economic Review*, 2014, 104 (6), 1518–1550.
- Crosen, Rachel and Uri Gneezy**, “Gender Differences in Preferences,” *Journal of Economic Literature*, 2009, 47 (2), 1365–1384.

- Diamond, Peter and Emmanuel Saez**, “The Case for a Progressive Tax: from Basic Research to Policy Recommendations,” *The Journal of Economic Perspectives*, 2011, pp. 165–190.
- Durante, Ruben, Louis Putterman, and Joël van der Weele**, “Preferences for Redistribution and Perception of Fairness: An Experimental Study,” *Journal of the European Economic Association*, forthcoming.
- Fehr, Ernst and Klaus M. Schmidt**, “A Theory Of Fairness, Competition, and Cooperation,” *Quarterly Journal of Economics*, 1999, *114* (3), 817–868.
- Fisman, Raymond, Pamela Jakiela, and Shachar Kariv**, “How Did Distributional Preferences Change During the Great Recession?,” *Journal of Public Economics*, 2015, *128*, 84–95.
- , — , — , and **Daniel Markovits**, “The Distributional Preferences of an Elite,” *Science*, 2015, *349* (6254), aab0096.
- , **Shachar Kariv, and Daniel Markovits**, “Individual Preferences for Giving,” *American Economic Review*, 2007, *97* (5), 1858–1876.
- Forsythe, Robert, Joel Horowitz, N. S. Savin, and Martin Sefton**, “Fairness in Simple Bargaining Games,” *Games and Economic Behavior*, 1994, *6* (3), 347–369.
- Harsanyi, John C**, “Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility,” *Journal of Political Economy*, 1955, *63*, 309–321.
- , “Bayesian Decision Theory and Utilitarian Ethics,” *American Economic Review*, 1978, *68* (2), 223–228.
- Hayes, Thomas**, “American Attitudes Toward Economic Inequality,” working paper 2011.

Henrich, Joseph, Jean Ensminger, Richard McElreath, Abigail Barr, Clark Barrett, Alexander Bolyanatz, Juan Camilo Cardenas, Michael Gurven, Edwins Gwako, Natalie Henrich, Carolyn Lesorogol, Frank Marlowe, David Tracer, and John Ziker, “Markets, Religion, Community Size, and the Evolution of Fairness and Punishment,” *Science*, 2010, *327* (5972), 1480–1484.

—, **Richard McElreath, Abigail Barr, Jean Ensminger, Clark Barrett, Alexander Bolyanatz, Juan Camilo Cardenas, Michael Gurven, Edwins Gwako, Natalie Henrich, Carolyn Lesorogol, Frank Marlowe, David Tracer, and John Ziker**, “Costly Punishment Across Human Societies,” *Science*, 2006, *312* (5781), 1767–1770.

Hermann, Benedikt, Christian Thöni, and Simon Gächter, “Antisocial Punishment Across Societies,” *Science*, 2008, *319* (5868), 1362–1367.

Kuziemko, Ilyana, “Redistributive preferences in an age of rising inequality,” unpublished presentation, in progress.

—, **Michael I Norton, Emmanuel Saez, and Stefanie Stantcheva**, “How Elastic Are Preferences for Redistribution? Evidence from Randomized Survey Experiments,” NBER Working Paper 18865, National Bureau of Economic Research 2013.

Lee, Woojin and John E Roemer, “Racism and Redistribution in the United States: A Solution to the Problem of American Exceptionalism,” *Journal of Public Economics*, 2006, *90* (6), 1027–1052.

Levine, David K., “Modeling Altruism and Spitefulness in Experiments,” *Review of Economic Dynamics*, 1998, *1* (3), 593–622.

- Loewenstein, George F, Leigh Thompson, and Max H Bazerman**, “Social Utility and Decision Making in Interpersonal Contexts,” *Journal of Personality and Social psychology*, 1989, *57* (3), 426.
- Meltzer, Allan H and Scott F Richard**, “A Rational Theory of the Size of Government,” *The Journal of Political Economy*, 1981, pp. 914–927.
- Mirrlees, JA**, “An Exploration in the Theory of Optimum Income Taxation,” *The Review of Economic Studies*, 1971, *38* (2), 175–208.
- Norton, Michael I and Dan Ariely**, “Building a Better America—One Wealth Quintile at a Time,” *Perspectives on Psychological Science*, 2011, *6* (1), 9–12.
- Rabin, Matthew**, “Incorporating Fairness into Game Theory and Economics,” *The American economic review*, 1993, pp. 1281–1302.
- Rawls, John**, *A Theory of Justice*, Cambridge, MA: Harvard University Press, 1971.
- Roth, Alvin E., Vesna Prasnikar, Masahiro Okuno-Fujiwara, and Shmuel Zamir**, “Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study,” *American Economic Review*, 1991, *81* (5), 1068–1095.
- Saez, Emmanuel and Stefanie Stantcheva**, “Generalized Social Marginal Welfare Weights for Optimal Tax Theory,” Technical Report 18835, National Bureau of Economic Research 2013.
- Varian, Hal R.**, “The Nonparametric Approach to Demand Analysis,” *Econometrica*, 1982, *50* (4), 945–972.
- , “Non-Parametric Tests of Consumer Behaviour,” *Review of Economic Studies*, 1983, *50* (1), 99–110.

Table 1: Classifying Subjects' Distributional Preference

	FAIR-MINDED	INTERMEDIATE	SELFISH	ALL SUBJECTS
EQUALITY-FOCUSED	12.2	35.7	10.5	58.4
EFFICIENCY-FOCUSED	18.5	17.6	5.6	41.6
ALL SUBJECTS	30.6	53.3	16.1	100.0

The numbers indicate the percentage of subjects in each cell. We classify a subject as fair-minded if $0.45 < \hat{\alpha}_n < 0.55$; a subject is classified as selfish if $\hat{\alpha}_n > 0.95$. We classify a subject as equality-focused (resp. efficiency-focused) if $\hat{\rho}_n < 0$ (resp. $\hat{\rho}_n > 0$). We obtain similar results using statistical tests to classify subjects' types.

Table 2: Comparing ALP Experimental Subjects with the US Population

	COMPLETED EXPERIMENT (1)	STARTED EXPERIMENT (2)	INVITED TO EXPERIMENT (3)	ENTIRE ALP (4)	US ADULTS (5)
Female	0.58	0.59	0.60	0.60	0.51
Age	49.37	49.71	48.41	49.05	46.68
18 to 44 years old	0.38	0.37	0.42	0.41	0.48
At least 65 years old	0.17	0.18	0.16	0.18	0.18
Caucasian (including Hispanics)	0.77	0.76	0.75	0.74	0.76
African American	0.11	0.12	0.12	0.12	0.12
Native American	0.01	0.01	0.01	0.01	0.01
Asian or Pacific Islander	0.02	0.02	0.02	0.02	0.05
Hispanic or Latino	0.18	0.19	0.19	0.21	0.15
High school diploma	0.91	0.91	0.91	0.93	0.88
College degree	0.31	0.29	0.30	0.36	0.27
Currently employed	0.56	0.54	0.58	0.58	0.59
Currently unemployed	0.11	0.11	0.11	0.10	0.06
Out of labor force	0.34	0.34	0.32	0.32	0.35
Lives in northeast (census region I)	0.18	0.19	0.19	0.17	0.18
Lives in midwest (census region II)	0.20	0.19	0.18	0.19	0.21
Lives in south (census region III)	0.35	0.34	0.34	0.34	0.37
Lives in west (census region IV)	0.27	0.27	0.29	0.29	0.23

Column 1 includes data on all ALP respondents who completed the experiment. Column 2 includes data on all ALP respondents who completed the experiment plus those who logged in to the experiment but did not complete it. Column 3 includes data on all ALP respondents who completed the experiment, those ALP respondents who logged in to the experiment but did not complete it, and those ALP respondents who were invited to participate in the experiment but never logged in to the website. Column 4 includes data on all ALP respondents. Column 5 includes data from the American Community Survey administered by the Census Bureau. The ACS interviewed more than 2.4 million respondents in 2012. Averages are weighted to reflect the adult population of the United States.

Table 3: OLS Regressions of Estimated $\hat{\alpha}_n$ Parameters on Subject Characteristics

<i>Dependent Variable:</i>	ESTIMATED $\hat{\alpha}_n$	
	(1)	(2)
Female	-0.015 (0.014)	-0.021 (0.014)
Youngest quartile (age 37 or less)	-0.003 (0.016)	-0.004 (0.017)
Oldest quartile (over 60)	0.026 (0.018)	0.025 (0.018)
Did not complete high school	-0.046** (0.019)	-0.039** (0.02)
Completed college	0.009 (0.016)	0.007 (0.016)
African American	-0.063*** (0.019)	-0.066*** (0.02)
Hispanic/Latino	-0.018 (0.017)	-0.017 (0.019)
Lowest income quartile	-0.0004 (0.017)	-0.002 (0.017)
Highest income quartile	-0.002 (0.018)	-0.002 (0.018)
Employed	0.003 (0.017)	0.005 (0.017)
Unemployed	-0.026 (0.023)	-0.03 (0.023)
Married	0.002 (0.019)	-0.004 (0.019)
Widowed, separated, or divorced	-0.016 (0.021)	-0.011 (0.022)
Catholic	-0.029 (0.018)	-0.038* (0.019)
Protestant	0.006 (0.018)	-0.003 (0.019)
No religious preference	-0.018 (0.018)	-0.016 (0.018)
Constant	0.704*** (0.029)	0.714*** (0.03)
State of Residence FEs	No	Yes
Observations	1002	1002
R^2	0.041	0.089

Robust standard errors in parentheses. All regressions include controls for respondents who are missing data on race (2), household income (5), or religious affiliation (8).

Table 4: OLS Regressions of Estimated $\hat{\rho}_n$ Parameters on Subject Characteristics

<i>Dependent Variable:</i> <i>Specification:</i> <i>Subjects included:</i>	ESTIMATED $\hat{\rho}_n$		ESTIMATED $\hat{\rho}_n$		ESTIMATED $\hat{\rho}_n$		DECILE OF $\hat{\rho}_n$		ρ_{high}	
	OLS REGRESSION		OLS REGRESSION		MEDIAN REGRESSION		OLS REGRESSION		OLS REGRESSION	
	ALL	NON-SELFISH	ALL	NON-SELFISH	ALL	NON-SELFISH	ALL	NON-SELFISH	ALL	NON-SELFISH
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Female	-0.94** (0.396)	-0.876** (0.415)	-1.062*** (0.402)	-0.996** (0.428)	-0.137** (0.06)	-0.117* (0.063)	-0.497*** (0.191)	-0.381** (0.191)	-0.066** (0.032)	-0.048 (0.033)
Youngest quartile (age 37 or less)	1.418*** (0.414)	1.387*** (0.428)	1.457*** (0.43)	1.458*** (0.449)	0.163** (0.074)	0.149* (0.078)	0.675*** (0.226)	0.594*** (0.224)	0.088** (0.04)	0.081** (0.041)
Oldest quartile (over 60)	0.017 (0.599)	0.087 (0.621)	-0.096 (0.603)	-0.022 (0.629)	-0.095 (0.081)	-0.085 (0.085)	-0.237 (0.262)	-0.199 (0.261)	-0.06 (0.044)	-0.055 (0.044)
Did not complete high school	0.057 (0.673)	0.119 (0.678)	0.417 (0.684)	0.488 (0.693)	0.123 (0.107)	0.133 (0.11)	0.38 (0.326)	0.498 (0.326)	0.101* (0.058)	0.118** (0.059)
Completed college	-0.096 (0.469)	-0.119 (0.49)	-0.375 (0.485)	-0.364 (0.508)	0.144** (0.07)	0.147** (0.074)	0.295 (0.22)	0.278 (0.218)	0.046 (0.037)	0.047 (0.038)
African American	0.747 (0.672)	0.898 (0.683)	0.412 (0.722)	0.574 (0.732)	0.313*** (0.1)	0.339*** (0.103)	0.657** (0.32)	0.827*** (0.321)	0.151*** (0.055)	0.176*** (0.056)
Hispanic/Latino	0.111 (0.551)	0.204 (0.566)	0.171 (0.624)	0.315 (0.648)	0.042 (0.084)	0.062 (0.088)	-0.086 (0.266)	-0.01 (0.263)	0.019 (0.046)	0.031 (0.046)
Lowest income quartile	1.137** (0.512)	1.220** (0.529)	1.041** (0.525)	1.077** (0.541)	0.259*** (0.077)	0.239*** (0.08)	0.484** (0.238)	0.525** (0.238)	0.078* (0.042)	0.084** (0.042)
Highest income quartile	-0.622 (0.533)	-0.658 (0.564)	-0.722 (0.532)	-0.745 (0.562)	0.005 (0.08)	-0.039 (0.085)	-0.013 (0.253)	-0.003 (0.253)	0.031 (0.043)	0.036 (0.044)
Employed	0.918* (0.543)	0.93 (0.567)	0.651 (0.53)	0.626 (0.555)	0.128* (0.074)	0.104 (0.078)	0.025 (0.244)	0.006 (0.242)	-0.009 (0.04)	-0.015 (0.04)
Unemployed	0.372 (0.725)	0.375 (0.747)	0.003 (0.731)	-0.05 (0.758)	0.044 (0.109)	0.04 (0.114)	-0.015 (0.332)	-0.002 (0.329)	0.003 (0.058)	0.005 (0.059)
Married	-0.013 (0.503)	0.023 (0.527)	-0.006 (0.522)	0.027 (0.549)	0.011 (0.086)	0.006 (0.091)	-0.173 (0.267)	-0.112 (0.262)	-0.028 (0.047)	-0.021 (0.048)
Widowed, separated, or divorced	-0.644 (0.626)	-0.613 (0.641)	-0.33 (0.658)	-0.303 (0.676)	-0.111 (0.099)	-0.106 (0.104)	-0.335 (0.311)	-0.251 (0.307)	-0.039 (0.055)	-0.027 (0.055)
Catholic	0.856 (0.58)	0.996* (0.601)	0.603 (0.583)	0.738 (0.607)	0.107 (0.084)	0.153* (0.088)	0.46* (0.274)	0.655** (0.273)	0.056 (0.046)	0.082* (0.046)
Protestant	0.283 (0.576)	0.349 (0.604)	0.369 (0.595)	0.433 (0.626)	-0.126 (0.082)	-0.139 (0.086)	-0.166 (0.259)	-0.076 (0.259)	-0.012 (0.044)	-0.006 (0.045)
No religious preference	-0.132 (0.542)	-0.034 (0.565)	-0.217 (0.555)	-0.154 (0.58)	-0.135* (0.081)	-0.118 (0.085)	-0.119 (0.259)	-0.02 (0.258)	-0.012 (0.044)	-0.00009 (0.045)
Constant	-3.408*** (0.841)	-3.749*** (0.889)	-3.013*** (0.834)	-3.334*** (0.889)	-0.382*** (0.127)	-0.409*** (0.135)	5.455*** (0.408)	5.045*** (0.4)	0.394*** (0.069)	0.336*** (0.07)
State of Residence FEs	No	No	Yes	Yes	No	No	No	No	No	No
Observations	1002	957	1002	957	1002	957	1002	957	1002	957
R^2	0.044	0.044	0.107	0.108	.	.	0.053	0.056	0.049	0.055

Robust standard errors in parentheses. All regressions include controls for respondents who are missing data on race (2), household income (5), or religious affiliation (8). NON-SELFISH subjects are those who allocated themselves no more than 99 percent of the tokens, on average. ρ_{high} is an indicator for being efficiency-focused in the sense of having an estimated $\hat{\rho}_n$ of at least 0.

Table 5: OLS Regressions of Likelihood of Voting for Obama in 2012 (on $\hat{\rho}_n$)

	— WITHOUT STATE FES —			— WITH STATE FES —		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: All Subjects</i>						
$\hat{\rho}_n$	-0.003 (0.003)	.	.	-0.005* (0.003)	.	.
Decile of $\hat{\rho}_n$.	-0.009 (0.006)	.	.	-0.013** (0.006)	.
ρ_{high} (i.e. $\hat{\rho}_n \geq 0$)	.	.	-0.045 (0.033)	.	.	-0.068** (0.034)
Demographic Controls	Yes	Yes	Yes	Yes	Yes	Yes
State of Residence FEs	No	No	No	Yes	Yes	Yes
Observations	766	766	766	766	766	766
<i>Panel B: Non-Selfish Subjects</i>						
$\hat{\rho}_n$	-0.004 (0.003)	.	.	-0.006* (0.003)	.	.
Decile of $\hat{\rho}_n$.	-0.012** (0.006)	.	.	-0.016** (0.006)	.
ρ_{high} (i.e. $\hat{\rho}_n \geq 0$)	.	.	-0.057* (0.034)	.	.	-0.077** (0.035)
Demographic Controls	Yes	Yes	Yes	Yes	Yes	Yes
State of Residence FEs	No	No	No	Yes	Yes	Yes
Observations	734	734	734	734	734	734

Robust standard errors in parentheses. All regressions include controls for respondents who are missing data on race (2), household income (5), or religious affiliation (8).

Table 6: OLS Regressions of Likelihood of Being a Democrat (on $\hat{\rho}_n$)

	— WITHOUT STATE FES —			— WITH STATE FES —		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: All Subjects</i>						
$\hat{\rho}_n$	-0.002 (0.003)	.	.	-0.005 (0.003)	.	.
Decile of $\hat{\rho}_n$.	-0.014* (0.007)	.	.	-0.02*** (0.007)	.
ρ_{high} (i.e. $\hat{\rho}_n \geq 0$)	.	.	-0.075* (0.04)	.	.	-0.104** (0.042)
Demographic Controls	Yes	Yes	Yes	Yes	Yes	Yes
State of Residence FEs	No	No	No	Yes	Yes	Yes
Observations	528	528	528	528	528	528
<i>Panel B: Non-Selfish Subjects</i>						
$\hat{\rho}_n$	-0.003 (0.003)	.	.	-0.005 (0.003)	.	.
Decile of $\hat{\rho}_n$.	-0.016** (0.007)	.	.	-0.023*** (0.008)	.
ρ_{high} (i.e. $\hat{\rho}_n \geq 0$)	.	.	-0.087** (0.041)	.	.	-0.112** (0.044)
Demographic Controls	Yes	Yes	Yes	Yes	Yes	Yes
State of Residence FEs	No	No	No	Yes	Yes	Yes
Observations	505	505	505	505	505	505

Robust standard errors in parentheses. All regressions include controls for respondents who are missing data on race (2), household income (5), or religious affiliation (8).

Table 7: OLS Regressions of Likelihood of Voting for Obama in 2012 (on $\hat{\alpha}_n$)

	— WITHOUT STATE FES —				— WITH STATE FES —			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Panel A: All Subjects</i>								
$\hat{\alpha}_n$	0.048 (0.086)	.	.	.	0.062 (0.086)	.	.	.
Decile of $\hat{\alpha}_n$.	0.004 (0.006)	.	.	.	0.004 (0.006)	.	.
$\hat{\alpha}_{high}$ (i.e. above median $\hat{\alpha}_n$)	.	.	0.012 (0.04)	0.016 (0.049)	.	.	0.023 (0.04)	0.02 (0.05)
$\hat{\alpha}_{high} \times$ lowest income quartile	.	.	.	-0.146 (0.183)	.	.	.	-0.058 (0.178)
$\hat{\alpha}_{high} \times$ highest income quartile	.	.	.	0.076 (0.183)	.	.	.	0.077 (0.189)
Demographic Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State of Residence FEs	No	No	No	No	Yes	Yes	Yes	Yes
Observations	766	766	766	766	766	766	766	766
<i>Panel B: Non-Selfish Subjects</i>								
$\hat{\alpha}_n$	-0.008 (0.089)	.	.	.	0.032 (0.09)	.	.	.
Decile of $\hat{\alpha}_n$.	-0.0002 (0.006)	.	.	.	0.002 (0.006)	.	.
$\hat{\alpha}_{high}$ (i.e. above median $\hat{\alpha}_n$)	.	.	-0.012 (0.042)	-0.0003 (0.051)	.	.	0.011 (0.043)	0.012 (0.052)
$\hat{\alpha}_{high} \times$ lowest income quartile	.	.	.	-0.176 (0.188)	.	.	.	-0.075 (0.183)
$\hat{\alpha}_{high} \times$ highest income quartile	.	.	.	0.049 (0.193)	.	.	.	0.058 (0.197)
Demographic Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State of Residence FEs	No	No	No	No	Yes	Yes	Yes	Yes
Observations	734	734	734	734	734	734	734	734

Robust standard errors in parentheses. All regressions include controls for respondents who are missing data on race (2), household income (5), or religious affiliation (8).

Table 8: OLS Regressions of Likelihood of Being a Democrat (on $\hat{\alpha}_n$)

	— WITHOUT STATE FES —			— WITH STATE FES —				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Panel A: All Subjects</i>								
$\hat{\alpha}_n$	0.027 (0.107)	.	.	.	0.044 (0.11)	.	.	.
Decile of $\hat{\alpha}_n$.	0.003 (0.007)	.	.	.	0.006 (0.008)	.	.
$\hat{\alpha}_{high}$ (i.e. above median $\hat{\alpha}_n$)	.	.	0.002 (0.049)	0.006 (0.059)	.	.	0.016 (0.051)	0.018 (0.061)
$\hat{\alpha}_{high} \times$ lowest income quartile	.	.	.	-0.168 (0.237)	.	.	.	-0.076 (0.247)
$\hat{\alpha}_{high} \times$ highest income quartile	.	.	.	0.085 (0.212)	.	.	.	0.035 (0.218)
Demographic Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State of Residence FEs	No	No	No	No	Yes	Yes	Yes	Yes
Observations	528	528	528	528	528	528	528	528
<i>Panel B: Non-Selfish Subjects</i>								
$\hat{\alpha}_n$	-0.03 (0.113)	.	.	.	0.002 (0.115)	.	.	.
Decile of $\hat{\alpha}_n$.	-0.0001 (0.008)	.	.	.	0.003 (0.008)	.	.
$\hat{\alpha}_{high}$ (i.e. above median $\hat{\alpha}_n$)	.	.	-0.03 (0.052)	-0.02 (0.062)	.	.	-0.011 (0.054)	-0.01 (0.064)
$\hat{\alpha}_{high} \times$ lowest income quartile	.	.	.	-0.125 (0.245)	.	.	.	-0.001 (0.257)
$\hat{\alpha}_{high} \times$ highest income quartile	.	.	.	0.02 (0.225)	.	.	.	-0.001 (0.228)
Demographic Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State of Residence FEs	No	No	No	No	Yes	Yes	Yes	Yes
Observations	505	505	505	505	505	505	505	505

Robust standard errors in parentheses. All regressions include controls for respondents who are missing data on race (2), household income (5), or religious affiliation (8).

Figure 1: Prototypical Fair-minded Distributional Preferences






Figure 2: Estimated $\hat{\alpha}_n$ Parameters, by Sub-Group




Dots indicate mean values. Circles indicate 25th and 75th percentiles. Bars indicate 95 percent confidence intervals for means.

Figure 3: Estimated Mean $\hat{\rho}_n$ Parameters, by Sub-Group



Dots indicate mean values. Bars indicate 95 percent confidence intervals for means.

Figure 4: Estimated Median $\hat{\rho}_n$ Parameters, by Sub-Group



Dots indicate median values. Circles indicate 25th and 75th percentiles. Bars indicate 95 percent confidence intervals for medians.