

The Distributional Preferences of Americans*

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

August 28, 2014

Abstract

This paper characterizes the distributional preferences of a large and diverse sample of Americans. We decompose distributional preferences into two distinct components — fair-mindedness (tradeoffs between oneself and others) and equality-efficiency tradeoffs — by estimating constant elasticity of substitution utility functions at the individual level. 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 demographic and socioeconomic characteristics. After controlling for individual characteristics, our experimental measure of equality-efficiency tradeoffs predicts the political decisions of our subjects; there is no significant relationship between our experimental measure of fair-mindedness and political preferences. Since optimal tax policy depends on the distributional preferences of voters and taxpayers, our findings can inform debates about government redistribution.

JEL Classification Numbers: C91, D64

*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 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: rf250@columbia.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, and the weight on reducing differences in incomes versus increasing total income. In 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. However, as a comparison of Harsanyi (1955) and Rawls (1971) would suggest, fair-minded people may disagree about the extent to which efficiency should be sacrificed to combat inequality.¹

Political debates often center on the redistribution of income. 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 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

¹Harsanyi (1955) and Rawls (1971) 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).

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.

To this end, we conduct an incentivized experiment using the American Life Panel (ALP), a longitudinal survey administered online by the RAND Corporation. The sample consists of more than 4,500 U.S. households and 5,000 individual members 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. The ALP thus provides an uncommon opportunity to bring together rich experimental and survey data on a diverse set of participants to study the distributional preferences of the U.S. population.

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 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

²The experimental design and data analysis draw on our prior work. The 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. This allows us to analyze behavior at the level of the individual subject, without the need to pool data or assume that subjects are homogenous.

objective.

The consistency of individual decisions naturally leads us to ask what kind of distributional preferences are consistent with the observed choices. In 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.

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. Only 2 displayed utilitarian preferences, while 3 displayed Cobb-Douglas preferences with equal weights on *self* and *other*. 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 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 redistribution context, 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 (i.e. the relative weight on *self* versus *other*) and ρ characterizes equality-efficiency tradeoffs (i.e.

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 estimation is done for each subject n separately, generating individual-level estimates $\hat{\alpha}_n$ and $\hat{\rho}_n$.

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, intermediate, or efficiency-focused. For 330 subjects (32.9 percent) we cannot reject the null hypothesis of fair-mindedness ($\hat{\alpha}_n = 1/2$); while we cannot reject the hypothesis of selfishness ($\hat{\alpha}_n = 1$) for 151 subjects (15.1 percent). Thus, fair-minded subjects outnumber selfish ones by more than two to one. More than half of our subjects display a statistically significant degree of either equality or efficiency focus, with the 285 efficiency-focused subjects (28.4 percent) slightly outnumbering the 245 subjects (24.4 percent) who are equality-focused. We observe a greater degree of efficiency-focus among fair-minded subjects: 63 (19.0 percent of fair-minded subjects) show a preference for equalizing payoffs while 120 (36.3 percent) show a preference for maximizing the average payoff.

Exploiting the detailed demographic and economic data available on ALP subjects, 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 great 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.

Given the potential link from distributional preferences to policy choices, it is natural to examine the empirical relationship between distributional preferences and subjects' political decisions. Though Democratic voters are typically assumed to be more equality-focused than Republicans, the recent work of Kuziemko, Norton, Saez, and Stantcheva (2013) highlights the possibility that supporters of the two parties may differ in their attitudes toward government intervention in general, and not in terms of their aversion to income inequality and their overall support for redistribution (once government is taken out of the equation). In our final set of results, we use data from ALP modules on voting and party affiliations to assess the relation-

ship between our experimental measures of distributional preferences and subjects' political choices, in a manner that can distinguish this relationship from attitudes toward government in general.

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 trade-offs.

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 (high) income households are more likely to affiliate with the Democrats (Republicans). However, because 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.

Overall, our findings contribute to the discussion around tax policy and other forms of government redistribution. In a standard model of taxation (Mirrlees, 1972), 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.³ 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 also 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. (i.e., another ALP respondent). This stands in

³See, for example, Benabou and Ok (2001) for the role of upward mobility; Lee and Romer (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.

contrast to many experiments, where subjects are generally matched with someone from their own community.

2 Related Literature

Experimental research has been very fruitful in both establishing the empirical reliability of distributional preferences and directing theoretical attention to 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, though 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.

We build most directly on our own prior work: Fisman, Kariv, and Markovits (2007) extends the modified dictator game first first proposed by Andreoni and Miller (2002); in our earlier paper, we introduced an experimental technique that enables us to collect richer individual-level data than had previously been possible. This is particularly important given that, as Andreoni and Miller (2002) emphasize, individual heterogeneity requires behavior to be examined at an individual level for distributional preferences to be properly understood.

Below, we focus attention on studies conducted with larger and more diverse pools of subjects than the usual collection of university students, as these are particularly relevant to our research. 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.⁴

As a result, there are several key differences between our study and that of Bellemare, Kröger, and van Soest (2008), beyond the fact that we draw our samples from different societies. First, 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. Second, their study explores 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 overall findings resonate with theirs. In particular, they also find considerable heterogeneity in preferences, much of which is not correlated with observable characteristics. Finally, building on our ability to measure efficiency-equality tradeoffs in the U.S. population, our findings contribute to central policy questions on the link from distributional preferences to voting decisions.

Like Bellemare, Kröger, and van Soest (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

⁴While Bellemare, Kröger, and van Soest (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, Kröger, and van Soest (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.

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.⁵ To recruit subjects for our experiment, ALP administrators sent email invitations to a random sample of the more than 4,000 ALP respondents for whom detailed demographic information is available. 1,172 received the email and logged in to the experiment.⁶ 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 (the most recent data available from the ACS). 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 geographic diversity; the subsample of 1,002 ALP respondents that constitute our subject pool is remarkably consistent with the entire ALP sample.⁷

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

⁵The 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.

⁶We 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).

⁷In the Online Appendix, we examine the individual characteristics associated with completing the experiment. The overall completion rate is high but there is some heterogeneity in across the socioeconomic and demographic characteristics of our subjects.

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. In our prior work, we developed a computer interface for exactly this purpose (Fisman, Kariv, and Markovits 2007). 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.⁸

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*.⁹ Confronting subjects with a rich menu of such budget lines

⁸The experimental method is applicable to many types of individual choice problems. See Choi, Fisman, Gale, and Kariv (2007) and Ahn, Choi, Gale, and Kariv (forthcoming), for settings involving, respectively, risk and ambiguity. ? 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.

⁹See Fisman, Kariv, and Markovits (2007) for an extended description of the experimental interface. Full experimental instructions are included in the Online Appendix. In a standard dictator experiment (cf. Forsythe, Horowitz, Savin, and Sefton 1994), $\pi_s + \pi_o = 1$:

allows us to identify both the tradeoff between both *self* and *other* (i.e. fair-mindedness) and the tradeoff between equality and efficiency.

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 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 . 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.

We conclude this section by reiterating that our design is particularly well-suited to assessing voters' attitudes toward redistribution—one of the main focuses of our analysis—as subjects make tradeoffs between their own payoff and the payoff an individual drawn from the general population of the U.S. (i.e., another ALP respondent), whereas in most experiments subjects are matched with someone from their own community.”

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), classical 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

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.

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 conclude 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 $\rho \leq 1$.¹⁰ 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.¹¹ Among fair-minded subjects, the family of

¹⁰The 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.

¹¹For $\alpha = 1/2$ and any ρ , (π, π') is chosen subject to the budget constraint $p_s\pi_s + p_o\pi_o = 1$ if and only if (π', π) is chosen subject to the mirror-image budget constraint $p_o\pi_s + p_s\pi_o = 1$.

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.

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.¹² We generate individual-level estimates \hat{g}_n and \hat{r}_n using non-linear Tobit maximum likelihood, and use these estimates to infer the values of the underlying CES parameters $\hat{\alpha}_n$ and $\hat{\rho}_n$.

¹²For 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.

4.2 Fair-mindedness

We classify a subject as fair-minded if we cannot reject the hypothesis that $\hat{\alpha}_n$ is equal to 0.5 at the 10 percent level. By this criterion, 330 subjects (32.9 percent) are fair-minded.¹³ In contrast, only 151 subjects (15.1 percent) are selfish in the sense that we cannot reject the hypothesis that their estimated $\hat{\alpha}_n$ is different from 1 at the 10 percent level.¹⁴ By this measure, fair-minded subjects outnumber selfish ones by approximately 2 to 1.¹⁵

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.¹⁶ We also find that non-Hispanic whites are significantly less fair-minded than both African American and Hispanic subjects.

In Table 3, we explore the associations between fair-mindedness and individual characteristics in a regression framework.¹⁷ 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 of residence fixed effects. The indicators for being African American and having less than a high school education are both negative and significant with and without state fixed effects, indicating that these groups are, on the whole, more fair-minded than other Americans.

¹³This definition involves both the estimated $\hat{\alpha}_n$ and the standard error associated with that estimate. Among fair-minded subjects, estimated $\hat{\alpha}_n$ parameters range from 0.258 to 0.749. The vast majority of fair-minded subjects (93.6 percent) have estimated $\hat{\alpha}_n$ parameters between 0.4 and 0.6.

¹⁴Since $\hat{\alpha}_n$ cannot be greater than 1, we use a one-sided test.

¹⁵We omit from our totals one subject who appears both selfish and fair-minded because the standard error associated with the $\hat{\alpha}_n$ estimate is quite large. This subject made inconsistent choices, as indicated by a CCEI of 0.293.

¹⁶The mean $\hat{\alpha}_n$ is 0.67 among women and 0.69 among men (p-value 0.06). However, using the binary indicator based on a test of the hypothesis that $\hat{\alpha}_n$ equals 1/2, we find that women are 6.0 percentage points more likely to be fair-minded (p-value 0.05).

¹⁷We 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 relatively few subjects have very high estimated $\hat{\alpha}_n$ parameters).

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 $\hat{\alpha}_n$. Adding state fixed effects raises the amount of variation that is explained by observables to 8.9 percent.

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. When we classify subjects as focused on equality (or efficiency) based on a one-sided statistical test of the hypothesis that $\hat{\rho}_n$ is less than (greater than) 0, the opposite pattern emerges. 285 subjects (28.4 percent) display a statistically significant level of efficiency focus, while 245 (24.5 percent) display a statistically significant level of equality focus. As these numbers clearly demonstrate, the American population includes large numbers of people holding divergent views on the relative importance of minimizing inequality versus maximizing efficiency. For more than half of our subjects, we can reject the hypothesis that $\hat{\rho}_n = 0$, suggesting that the majority of Americans have strongly held, but divergent, views on the tradeoff between equality and efficiency.

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.

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.

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 Voting 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 sur-

vey 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, Norton, Saez, and Stantcheva (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, 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.¹⁸

We explore the link between equality-efficiency tradeoffs and political behavior by looking at voting decisions in the 2012 presidential election. Our main dependent variable is an indicator for voting for Democrat Barack Obama, a relatively pro-redistribution candidate, rather than Republican Mitt Romney.¹⁹ 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.²⁰ 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

¹⁸Durante, 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.

¹⁹As 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].”

²⁰Unfortunately, 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.

distinct from their distributional preferences.²¹ We employ a linear probability model with an indicator variable for having voted for Obama as the outcome.²² We report results for all three measures of equality-efficiency tradeoffs used in the preceding 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 marginally with the inclusion of ρ_{high} , indicating that some amount of the gender voting gap can be directly accounted for by distributional preferences.)

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

²¹Interestingly, 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.

²²Results of probit estimations are nearly identical.

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.²³ 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.

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.

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.²⁴ Interestingly, we do not find sup-

²³Results 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.

²⁴An 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

port 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.

6 Conclusion

In this paper, we take a first step in characterizing, via experiments administered through the American Life Panel, 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 unexpected 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.

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 critical 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, Norton, Saez, and Stantcheva 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. Moreover, our results show that lower income Americans are more focused on efficiency than

(2005) and the references cited therein.

other groups; while, in related work, Kuziemko (in progress) and Fisman, Jakiela, and Kariv (2013) 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.

Our experiments are only a first step in a much larger agenda. Theories of redistribution should evolve alongside empirical evidence, with more research in the spirit of Saez and Stantcheva (2013) to better understand the implications of observed distributional preferences for optimal policy. Our results show that lab experiments can make a positive contribution to this discourse, providing measures of the extent to which fair-minded voters are willing to sacrifice efficiency to combat inequality, and predicting their level of support for government policies designed to achieve that objective.

References

- AFRIAT, S. N. (1967): “The Construction of Utility Functions from Expenditure Data,” *International Economic Review*, 8(1), 67–77.
- (1972): “Efficiency Estimates of Production Functions,” *International Economic Review*, 8, 568–598.
- AHN, D., S. CHOI, D. GALE, AND S. KARIV (forthcoming): “Estimating Ambiguity Aversion in a Portfolio Choice Experiment,” *Quantitative Economics*.
- ALESINA, A., AND E. LA FERRARA (2005): “Preferences for Redistribution in the Land of Opportunities,” *Journal of Public Economics*, 89(5), 897–931.
- ANDREONI, J., AND J. MILLER (2002): “Giving According to GARP: An Experimental Test of the Consistency of Preferences for Altruism,” *Econometrica*, 70(2), 737–753.
- ANDREONI, J., AND L. VESTERLUND (2001): “Which Is the Fairer Sex? Gender Differences in Altruism,” *Quarterly Journal of Economics*, 116(1), 293–312.
- BELLEMARE, C., AND S. KRÖGER (2007): “On Representative Social Capital,” *European Economic Review*, 51(1), 183–202.
- BELLEMARE, C., S. KRÖGER, AND A. VAN SOEST (2008): “Measuring Inequity Aversion in a Heterogeneous Population Using Experimental Decisions and Subjective Probabilities,” *Econometrica*, 76(4), 815–839.
- (2011): “Preferences, Intentions, and Expectation Violations: A Large-Scale Experiment with a Representative Subject Pool,” *Journal of Economic Behavior and Organization*, 78(3), 349–365.
- BERG, J. E., J. DICKHAUT, AND K. MCCABE (1995): “Trust, reciprocity, and social history,” *Games and Economic Behavior*, 10, 122–142.
- BOLTON, G. E. (1991): “A Comparative Model of Bargaining: Theory and Evidence,” *The American Economic Review*, pp. 1096–1136.
- BOLTON, G. E., AND A. OCKENFELS (1998): “Strategy and Equity: an ERC-Analysis of the Güth–van Damme Game,” *Journal of Mathematical Psychology*, 42(2), 215–226.
- (2000): “ERC: A Theory of Equity, Reciprocity, and Competition,” *American Economic Review*, pp. 166–193.
- CAMERER, C. (2003): *Behavioral Game Theory*. Princeton University Press, Princeton, NJ.
- CHARNESS, G., AND M. RABIN (2002): “Understanding Social Preferences with Simple Tests,” *Quarterly Journal of Economics*, 117(3), 817–869.
- (2005): “Expressed Preferences and Behavior in Experimental Games,” *Games and Economic Behavior*, 53(2), 151–169.

- CHOI, S., R. FISMAN, D. GALE, AND S. KARIV (2007): “Consistency and Heterogeneity of Individual Behavior under Uncertainty,” *American Economic Review*, 97(5), 1921–1938.
- CROSON, R., AND U. GNEEZY (2009): “Gender Differences in Preferences,” *Journal of Economic Literature*, 47(2), 1365–1384.
- DIAMOND, P., AND E. SAEZ (2011): “The Case for a Progressive Tax: from Basic Research to Policy Recommendations,” *The Journal of Economic Perspectives*, pp. 165–190.
- DURANTE, R., L. PUTTERMAN, AND J. VAN DER WEELE (forthcoming): “Preferences for Redistribution and Perception of Fairness: An Experimental Study,” *Journal of the European Economic Association*.
- FEHR, E., AND K. M. SCHMIDT (1999): “A Theory Of Fairness, Competition, and Cooperation,” *Quarterly Journal of Economics*, 114(3), 817–868.
- FISMAN, R., P. JAKIELA, AND S. KARIV (2013): “How Did the Great Recession Impact Social Preferences?,” working paper,.
- FISMAN, R., S. KARIV, AND D. MARKOVITS (2007): “Individual Preferences for Giving,” *American Economic Review*, 97(5), 1858–1876.
- FORSYTHE, R., J. HOROWITZ, N. S. SAVIN, AND M. SEFTON (1994): “Fairness in Simple Bargaining Games,” *Games and Economic Behavior*, 6(3), 347–369.
- HARSANYI, J. C. (1955): “Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility,” *Journal of Political Economy*, 63, 309–321.
- (1978): “Bayesian Decision Theory and Utilitarian Ethics,” *American Economic Review*, 68(2), 223–228.
- HAYES, T. (2011): “American Attitudes Toward Economic Inequality,” working paper.
- HENRICH, J., J. ENSMINGER, R. MCELREATH, A. BARR, C. BARRETT, A. BOLYANATZ, J. C. CARDENAS, M. GURVEN, E. GWAKO, N. HENRICH, C. LESOROGOL, F. MARLOWE, D. TRACER, AND J. ZIKER (2010): “Markets, Religion, Community Size, and the Evolution of Fairness and Punishment,” *Science*, 327(5972), 1480–1484.
- HENRICH, J., R. MCELREATH, A. BARR, J. ENSMINGER, C. BARRETT, A. BOLYANATZ, J. C. CARDENAS, M. GURVEN, E. GWAKO, N. HENRICH, C. LESOROGOL, F. MARLOWE, D. TRACER, AND J. ZIKER (2006): “Costly Punishment Across Human Societies,” *Science*, 312(5781), 1767–1770.
- HERMANN, B., C. THÖNI, AND S. GÄCHTER (2008): “Antisocial Punishment Across Societies,” *Science*, 319(5868), 1362–1367.
- KUZIEMKO, I. (in progress): “Redistributive preferences in an age of rising inequality,” unpublished presentation,.

- KUZIEMKO, I., M. I. NORTON, E. SAEZ, AND S. STANTCHEVA (2013): “How Elastic Are Preferences for Redistribution? Evidence from Randomized Survey Experiments,” NBER Working Paper 18865, National Bureau of Economic Research.
- LEVINE, D. K. (1998): “Modeling Altruism and Spitefulness in Experiments,” *Review of Economic Dynamics*, 1(3), 593–622.
- LOEWENSTEIN, G. F., L. THOMPSON, AND M. H. BAZERMAN (1989): “Social Utility and Decision Making in Interpersonal Contexts,” *Journal of Personality and Social Psychology*, 57(3), 426.
- NORTON, M. I., AND D. ARIELY (2011): “Building a Better America—One Wealth Quintile at a Time,” *Perspectives on Psychological Science*, 6(1), 9–12.
- RABIN, M. (1993): “Incorporating Fairness into Game Theory and Economics,” *The American economic review*, pp. 1281–1302.
- RAWLS, J. (1971): *A Theory of Justice*. Harvard University Press, Cambridge, MA.
- ROTH, A. E., V. PRASNIKAR, M. OKUNO-FUJIWARA, AND S. ZAMIR (1991): “Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study,” *American Economic Review*, 81(5), 1068–1095.
- SAEZ, E., AND S. STANTCHEVA (2013): “Generalized Social Marginal Welfare Weights for Optimal Tax Theory,” Discussion Paper 18835, National Bureau of Economic Research.
- VARIAN, H. R. (1982): “The Nonparametric Approach to Demand Analysis,” *Econometrica*, 50(4), 945–972.
- (1983): “Non-Parametric Tests of Consumer Behaviour,” *Review of Economic Studies*, 50(1), 99–110.

Table 1: Distributional Preference Types

	FAIR-MINDED	INTERMEDIATE	SELFISH	ALL SUBJECTS
EQUALITY-FOCUSED	6.3	14.6	3.6	24.5
INTERMEDIATE	14.7	24.5	8.0	47.1
EFFICIENCY-FOCUSED	12.0	13.0	3.5	28.4
ALL SUBJECTS	32.9	52.0	15.1	100.0

The numbers indicate the percentage of subjects in each cell. We classify a subject as fair-minded if we cannot reject the null hypothesis that $\hat{\alpha}_n = 1/2$; similarly, a subject is classified as selfish if we cannot reject the null that $\hat{\alpha}_n = 1$ (both tests are at the 10 percent level, though the test for selfishness is one-sided since $\hat{\alpha}_n = 1$ at the boundary of the parameter space). One subject who had many revealed preference violations is classified as both fair-minded and selfish, and is therefore included in the intermediate category. We classify a subject as equality-focused or efficiency-focused if we can reject the hypothesis that $\hat{\rho}_n = 0$ at the 10 percent level using a one-sided test. When we can reject the null in favor of the alternative hypothesis that $\hat{\rho}_n$ is less (greater) than 0, we classify a subject as being focused on equality (efficiency).

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.