NBER WORKING PAPER SERIES

THE DISTRIBUTIONAL PREFERENCES OF AMERICANS

Raymond Fisman
Pamela Jakiela
Shachar Kariv

Working Paper 20145
http://www.nber.org/papers/w20145

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2014

We thank Daniel Markovits and Daniel Silverman for helpful discussions and comments. This paper has also benefited from suggestions by the participants of seminars at several universities and conferences. 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. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

The Distributional Preferences of Americans
Raymond Fisman, Pamela Jakiela, and Shachar Kariv
NBER Working Paper No. 20145
May 2014
JEL No. C91,D64

ABSTRACT

We measure the distributional preferences of a large, diverse sample of Americans by embedding modified dictator games that vary the relative price of redistribution in the American Life Panel. Subjects' choices are generally consistent with maximizing a (social) utility function. 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. Approximately equal numbers of Americans have equality-focused and efficiency-focused distributional preferences. After controlling for individual characteristics, our experimental measures of equality-efficiency tradeoffs predict the political decisions of our subjects.

Raymond Fisman
School of Business
Columbia University
622 Uris Hall
3022 Broadway
New York, NY 10027
and NBER
rf250@columbia.edu

Pamela Jakiela
Department of Agricultural
and Resource Economics
University of Maryland
2200 Symons Hall
College Park, MD 20742
pjakiela@gmail.com

Shachar Kariv
Department of Economics
University of California, Berkeley
Berkeley, CA 94720
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 is often 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

[1] Harsanyi (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).

the importance of correctly distinguishing fair-mindedness from preferences regarding 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 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.[2]

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.[3] This design allows us to decompose

[2]Bellemare, Kröger, and van Soest (2008) study distributional preferences in a large, diverse sample of Dutch adults. In their experiment, survey respondents from the CentERpanel participate in ultimatum games. Data characterizing their decisions within the experiment, their beliefs about the likelihood that specific ultimatum game offers would be accepted, and their individual characteristics is used to estimate a structural model of inequality aversion in the Dutch population. Aside from different societies, we restrict attention to dictator games and ignore the complications of strategic behavior and reciprocity motivations in response games in order to focus on behavior motivated by purely distributional preferences. (While 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.)

As a result, there are two key differences between our study and that of Bellemare, Kröger, and van Soest (2008). 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.

[3]The 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.

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).

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.

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. The 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 relationship between our experimental measures of distributional preferences and subjects' political choices. 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.

To summarize, our experiment characterizes the distributional preferences of a diverse set of Americans. We find that subjects are consistent in their choices, and show a high degree of heterogeneity in both the extent of fair-mindedness and willingness to trade off efficiency and equality. Yet in spite of this heterogeneity, a high fraction of Americans across a wide range of demographic and economic characteristics are fair-minded in the sense that their choices are not skewed in favor of their own payoffs. Further, among fair-minded individuals, we observe the full spectrum of equality-efficiency tradeoffs. These findings may inform debates about tax policy and other forms of government redistribution. 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.

2 Experimental Design

We embed an incentivized experiment in the American Life Panel (ALP), an internet survey administered by the RAND Corporation to approximately 6,000 adult Americans. The ALP allows us to implement a sophisticated experiment with a large and diverse subject pool, providing a unique opportunity to explore the heterogeneity in distributional preferences within the U.S. population.

2.1 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.[4]

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 subjects 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*.[5] Confronting subjects with a rich menu of such budget lines allows us to identify both the tradeoff between both *self* and *other* (i.e.

[4]The 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. Choi, Kariv, Müller, and Silverman (forthcoming) 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.

[5]See 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), $\pi_s + \pi_o = 1$: 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.

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.

2.2 Subject Pool

The ALP is a panel survey of approximately 6,000 American adults implemented online by the RAND Corporation.[6] To recruit subjects for our experiment, ALP administrators sent email invitations to a random sample of the more than 5,000 ALP respondents for whom detailed demographic information is available. 1,172 received the email and logged in to the experiment.[7] Table 2 describes the progression of ALP respondents who logged in through the instructions and the incentivized decision problems.[8] 1,002 respondents — 85.5 percent of those who logged in — completed the experiment; these subjects constitute our subject pool.

[6]ALP respondents have been recruited in several different ways. The 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.

[7]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).

[8]In the Online Appendix, we examine the individual characteristics associated with completing the experiment in the sample of 1,172 ALP respondents who ever logged in. The completion rate are significantly higher for college graduates and significantly lower for African Americans.

Summary statistics on our 1,002 experimental subjects are reported in Table 3. Though they are not a random sample of Americans, they are broadly representative of the geographic, demographic, and socioeconomic diversity of the U.S. population. Subjects in our experiment are drawn from 47 U.S. states, and range in age from 19 to 91. 58.3 percent are female. 9.0 percent of our subjects did not finish high school, while 12.6 percent hold graduate degrees. The average household income of subjects in our sample is approximately 55,000 US dollars. 55.6 percent of subjects are currently employed; the remainder include retirees (17.4 percent), the unemployed (10.3 percent), the disabled (8.1 percent), homemakers (6.3 percent), and others who are on medical leave or otherwise temporarily absent from the work force. 67.6 percent identify themselves as non-Hispanic whites, 17.9 percent as Hispanic or Latino, and 11.0 percent as African American. 17.9 percent live in the Northeast (census region I), 20.2 percent in the Midwest (census region II), 35.1 percent in the South (census region III), and 26.8 percent in the West (census region IV). Our subject pool is therefore substantially more diverse in terms of age, educational attainment, household income, occupational status, and place of residence than the samples of university students used in most lab experiments. Implementing our experiment through the ALP also allows us to draw on the wealth of information that has been collected through previous surveys, and to combine this information with our experimental data.

3 Individual Heterogeneity

In this section, we explore the allocation decisions of our experimental subjects in a simple framework that imposes minimal functional form assumptions on distributional preferences.

3.1 Fair-mindedness

We begin by constructing a simple, reduced form measure of fair-mindedness: the fraction of tokens kept by *self*, $\pi_s/(\pi_s + \pi_o)$, averaged across all 50 decision problems at the subject level. This measure is equal to one for perfectly selfish subjects; fair-minded subjects who put equal weight on *self* and *other* will keep approximately half of the total tokens, on average. We observe considerable heterogeneity across subjects. The individual-level average of $\pi_s/(\pi_s+\pi_o)$ ranges from 0.03 to 1, though the vast majority of subjects (84.6 percent) kept an average of at least half the tokens. Only 35 subjects (3.49 percent) kept an average of less than 45 percent of tokens. Thus, among subjects that kept less than half of the tokens, most appear to place nearly equal weight on the payoffs to *self* and *other*.[9]

Several key features of our data stand out. First, we observe very low numbers of selfish subjects who kept almost all of the tokens. The average of $\pi_s/(\pi_s+\pi_o)$ is at least 95 percent For only 81 subjects (8.1 percent). This relatively low number of selfish subjects contrasts with the large body of experiments with the the usual collection of undergraduate students. Overall, our subjects kept approximately 65 percent of the tokens. In the studies of standard split-the-pie dictator games reported in Camerer (2003), the typical mean allocations to *other* are about 80 percent. Second, a substantial fraction of subjects kept an average of approximately half the tokens. In fact, these fair-minded subjects far outnumber the selfish types: 370 subjects (36.9 percent) kept an average of 45 to 55 percent of the tokens. Moreover, the distribution of $\pi_s/(\pi_s+\pi_o)$ is quite smooth between 0.5 and 0.99, suggesting considerable heterogeneity in fair-mindedness among non-selfish subjects.

Figure 2 explores the extent to which this heterogeneity in fair-mindedness is explained by demographic and socioeconomic characteristics. Each section of the figure represents a partition of the subject pool into mutually exclusive categories — for example, men and women. The figure indicates the average across subjects of the individual-level average of $\pi_s/(\pi_s+\pi_o)$ within a category; the 95 percent confidence intervals for means, and the 25^{th} and 75^{th} percentiles of the distribution are labeled for each group. There are substantial differences in the average of $\pi_s/(\pi_s+\pi_o)$ across groups. Women keep a smaller fraction of the tokens than men. Surprisingly, $\pi_s/(\pi_s+\pi_o)$ increases with both household income and education level. A number of these differences are statistically significant.

In addition to the clear between-group differences, there is considerable heterogeneity within every category. For all the sub-groups included in the figure, the 25^{th} percentile of the distribution is between 0.5 and 0.52. This means that every demographic and socioeconomic category we consider includes non-negligible numbers of fair-minded subjects who treat *self* and *other* more or less symmetrically. The 75^{th} percentiles range from 0.62 to 0.83, and we observe relatively selfish subjects who keep an average of at

[9] This suggests that almost all of our subjects comprehended the tradeoff between *self* and *other* that they were making. Numerous experimental studies suggest that subjects rarely allocate more to *other* than to *self* in standard dictator games. Since our design includes random variation in the price of redistribution and subjects may respond to price variation in different ways, subjects who put equal weight on the payoffs to *self* and *other* may not allocate themselves exactly half of the tokens in our experiment.

least 95 percent of the tokens in every category. In most cases — for example, when we compare men to women or lower and higher income households — there is much more variation among subjects within a category than there is across category averages. Regression analysis confirms this: the complete set of dummy variables for demographic and socioeconomic categories explains 4.95 percent of the variation in average $\pi_s/(\pi_s + \pi_o)$.[10] Thus, most of the observed heterogeneity in fair-mindedness is not explained by demographic and socioeconomic factors.

3.2 Equality-Efficiency Tradeoffs

Subjects may also differ in their equality-efficiency tradeoffs, as discussed above. Of the fair-minded subjects, 85 subjects (8.5 percent) always made nearly equal allocations $\pi_s = \pi_o$ indicating Rawlsian preferences.[11] Only three subject allocated all their tokens to *self* when $p_s < p_o$ and to *other* when $p_s > p_o$ indicating utilitarian preferences, while two subjects made equal expenditures on *self* and *other* $p_s\pi_s = p_o\pi_o$ indicating Cobb-Douglas preferences. Thus, very few subjects made allocations that fit with fair-minded prototypical distributional preferences.

To explore the equality-efficiency tradeoffs of the remaining subjects, we regress the budget share spent on tokens kept ($p_s\pi_s$) on the log-price of redistribution ($p = p_s/p_o$) at the individual level. We classify a subject as efficiency-oriented if the OLS slope coefficient is greater than or equal to 0 because increasing $p_s\pi_s$ when p increases indicates distributional preferences weighted towards efficiency (increasing total payoffs), whereas decreasing $p_s\pi_s$ when p increases indicates distributional preferences weighted towards equality (reducing differences in payoffs).

In Figure 3, we explore the variation in the fraction of efficiency-oriented subjects across demographic and socioeconomic categories. Each section of the figure represents a partition of the subject pool into mutually exclusive categories, and we indicate the proportions and the 95 percent confidence intervals. We again observe considerable variation within and across subgroups. Specifically, less educated subjects (those with less than a high school diploma), minorities, younger subjects, the unemployed, and the

[10]Regression results are reported in the Online Appendix.

[11]Since humans implement their decisions with error, we classify subjects as being consistent with a prototypical model of distributional preferences if, on average, their choices deviate from those prescribed by that model by less than 0.02 (i.e. by no more than 2 percent of the tokens or budget). In the Online Appendix, we report the fraction of subjects behaving in a manner consistent with each of the prototypical types for a range of values for the maximum average deviation.

never married are more efficiency-focused than other groups; older Americans, retirees, non-Hispanic whites, and Protestants focus less on efficiency and more on equality. As in the case of fair-mindedness, most of the observed heterogeneity in equality-efficiency tradeoffs is also not explained by demographic and socioeconomic factors.

4 Individual Rationality

In this section, we discuss our revealed preference tests of individual rationality in detail. Readers less interested in the technical aspects of our tests may prefer to bypass this portion of the paper and proceed to Section 5.

The most basic question to ask about choice data is whether it is consistent with individual utility maximization. If participants choose allocations subject to standard budget constraints (as in our experiment), classical revealed preference theory provides a direct test. Afriat's (1967) theorem shows that choices in a finite collection of budget sets are consistent with maximizing a well-behaved (piecewise linear, continuous, increasing, and concave) utility function $u_s(\pi_s, \pi_o)$ if and only if they satisfy the Generalized Axiom of Revealed Preference (GARP). Hence, to assess whether our data are consistent with utility-maximizing behavior, we only need to check whether our data satisfy GARP, which requires that if $\pi = (\pi_s, \pi_o)$ is indirectly revealed preferred to π', then π' is not directly revealed strictly preferred ($\mathbf{p}' \cdot \pi \geq \mathbf{p}' \cdot \pi'$) to π.

Although testing conformity with GARP is conceptually straightforward, there is an obvious difficulty: GARP provides an exact test of utility maximization – either the data satisfy GARP or they do not. 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), which measures the fraction by which each budget constraint must be shifted in order to remove all violations of GARP. By definition, the CCEI is bounded between zero and one. The closer it is to one, the smaller the perturbation of the budget constraints required to remove all violations and thus the closer the data are to satisfying GARP and hence to perfect consistency with utility maximization. The difference between the CCEI and one can be interpreted as an upper bound on the fraction of income that a subject is wasting by making inconsistent choices.

There is no natural threshold for the CCEI for determining whether subjects are close enough to satisfying GARP that they can considered utility maximizers. To generate a benchmark against which to compare our CCEI scores, we follow Bronars (1987), which builds on Becker (1962), and compare the behavior of our actual subjects to the behavior of simulated subjects who randomize uniformly on each budget line. Such tests are frequently applied to experimental data. The power of Bronars's (1987) test is defined to be the probability that a randomizing subject violates GARP. Choi, Fisman, Gale, and Kariv (2007) show there is a very high probability that even random behavior will pass the GARP test if the number of individual decisions is sufficiently low, underscoring the need to collect choices in a wide range of budget sets in order to provide a stringent test of utility maximization. In a simulation of 25,000 subjects who randomize uniformly on each budget line when confronted with our sequence of 50 decision problems, all the simulated subjects had GARP violations, so the Bronars criterion attains its maximum value.

The Bronars (1987) test rules out the possibility that consistency is the accidental result of random behavior, but it is not sufficiently powerful to detect whether utility maximization is the correct model. To this end, Fisman, Kariv, and Markovits (2007) generate a sample of hypothetical subjects who implement a CES utility function with an idiosyncratic preference shock that has a logistic distribution

$$\Pr(\pi^*) = \frac{e^{\gamma \cdot u(\pi^*)}}{\int_{\mathbf{p}\cdot\pi=1} e^{\gamma \cdot u(\pi)}}$$

where the precision parameter γ reflects sensitivity to differences in utility – the choice becomes purely random as γ goes to zero (Bronars' test), whereas the probability of the allocation yielding the highest utility approaches one as γ goes to infinity. The results provide a clear benchmark of the extent to which subjects do worse than choosing consistently and the extent to which they do better than choosing randomly, and demonstrate that if utility maximization is not in fact the correct model, then our experiment is sufficiently powerful to detect it. We refer the interested reader to Fisman, Kariv, and Markovits (2007) Appendix III for more detail.[12]

The CCEI scores in the ALP sample averaged 0.862 over all subjects, which we interpret as confirmation that most subjects' choices are approximately consistent. In comparison, the mean CCEI score of a sample of

[12] Varian (1982, 1983) modified Afriat's (1967) results and describes efficient and general techniques for testing the extent to which choices satisfy GARP. We refer the interested reader to Choi, Fisman, Gale, and Kariv (2007) for more details on testing for consistency with GARP and other measures that have been proposed for measuring GARP violations. In practice, all these measures yield similar conclusions.

25,000 random subjects ($\gamma = 0$) who made 50 choices from randomly generated budget sets in the same way as our human subjects is only 0.60. 74.2 percent of actual subjects have CCEI scores above 0.80, while 10.2 percent of random subjects have scores that high. If we choose the 0.85 efficiency level as our critical value, 64.1 percent of our subjects have CCEI scores above this threshold, while 3.4 percent of the random subjects have CCEI scores above 0.85.

There is, however, marked heterogeneity in the CCEI scores within and across the demographic and economic groups. Subjects that completed college display greater levels of consistency than subjects with less education. The magnitudes imply that, on average, subjects without college degrees waste 2.6 percentage points more of their earnings by making inconsistent choices relative to college graduates.[13] We also find that men are more consistent than women, and that the choices of white and Hispanic subjects are more consistent with utility maximization than those of African Americans in our sample. Though all three differences are statistically significant, they are small in magnitude; the average CCEI is above 0.8 for all the demographic and socioeconomic categories we consider.

5 Decomposing Distributional Preferences

5.1 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:

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

[13] Regression analysis of the relationship between CCEI scores and demographic and economic characteristics is provided in the Online Appendix.

where $0 \leq \alpha \leq 1$ and $\rho \leq 1$.[14] 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.[15] 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.

More generally, different values of ρ give different degrees to which equality is valued over efficiency. Specifically, 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 towards 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

[14]The 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.

[15]For $\alpha = 1/2$ 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 $p_o\pi_s + p_s\pi_o = 1$.

individual-level econometric specification for each subject n:

$$p^i_{s,n}\pi^i_{s,n} = \frac{g_n}{(p^i_{s,n}/p^i_{o,n})^{r_n} + g_n} + \epsilon^i_n$$

where $i = 1, ..., 50$ indexes the decision round and ϵ^i_n is assumed to be distributed normally with mean zero and variance σ^2_n. 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. [16] 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 $\hat{\rho}_n$.

5.2 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.[17] 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.[18] By this measure, fair-minded subjects outnumber selfish ones by approximately 2 to 1.[19]

Figure 4 explores the extent to which heterogeneity in estimated $\hat{\alpha}_n$ parameters is explained by observable characteristics. We again 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; and the unemployed and, to a lesser extent, the disabled appear more fair-minded

[16] For 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.

[17] This 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.

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

[19] We 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.

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

In Table 4, we explore the associations between fair-mindedness and individual characteristics in a regression framework.[21] In Columns 1 through 8, we consider each set of demographic or economic categories in isolation. Results are consistent with the patterns observed in Figure 4: when we do not control for other factors, gender, age, marital status, education level, household income, employment status, and (to some extent) religion are all significant predictors of where one falls on the spectrum from fair-mindedness to selfishness. In Columns 9 and 10, we include all the individual characteristics in the same specification; in Column 10, 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.

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 8.8 percent.

5.3 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

[20]The 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).

[21]We 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).

holding divergent views on the relative importance of minimizing inequality vis-a-vis 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 5 and 6, we disaggregate the estimated $\hat{\rho}_n$ parameters by demographic and socioeconomic categories (Figure 5 presents means and Figure 6 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 Tables 5 and 6, we explore the associations between equality-efficiency tradeoffs and individual characteristics in a regression framework. Table 5 replicates the OLS specifications from Table 4 with $\hat{\rho}_n$ as the dependent variable. Given the skewed distribution of the estimated $\hat{\rho}_n$ parameters, Table 6 reports a number of alternative specifications: we report median regressions in Columns 3 and 4, regressions in which the outcome variable is the decile of the estimated $\hat{\rho}_n$ distribution in Columns 5 and 6, and regressions in which the dependent variable is an indicator for having $\hat{\rho}_n \geq 0$ (which we term ρ_{high}) in Columns 7 and 8. 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.

Several robust associations, already hinted at by the patterns in Figures 5 and 6, stand out. First, the youngest quartile of subjects are significantly more efficiency-focused than older individuals (in all specifications). The coefficient in Column 7, for example, suggests that subjects in the youngest quartile are 8.8 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, African American subjects are more efficiency-focused than non-Hispanic whites (the omitted category). Though the coefficient is only significant in 6 of the 8 specifications, the point estimates are extremely large, suggesting, for example, that African Americans are 15.1 percent more likely to have a $\hat{\rho}_n$ of at least 0. Finally, though the associations are not signification in all specifications, we also find that women are substantially more equality-focused than men.

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.

6 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.[22] A vast literature on the partisan preferences of Americans assumes that Democrats have a stronger preference 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, 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

[22]In the Online Appendix, we explore the relationship between fair-mindedness and political preferences. An extensive literature explores the extent to which voters support policies that are in their own perceived short-run and long-run economic interests (cf. Alesina and La Ferrara (2005) and the references cited therein). 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 lean toward the Democrats (Republicans). However, because our measure of household income provides only a rough indicator of who is likely to benefit from government redistribution, we do not view our results as evidence that self-interest plays no role in voting decisions.

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.[23]

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.[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 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 7, we present specifications which

[23]Durante, 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 allows them to distinguish between three motives for supporting income redistribution – own-income maximization (of those in lower income brackets), risk aversion, and distributional preferences. The subjects in the 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.

[24]As 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]."

[25]Unfortunately, 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.

[26]Interestingly, 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.

[27]Results of probit estimation are nearly identical.

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 coefficient on ρ_{high} in Column 8 is -0.068, implying a marginally larger impact of distributional preferences on voting outcomes when geographic variation is accounted for. In Panel B, 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 aligning 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 8. 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 at least 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. Thus, 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 about government. However, this pattern only emerges after one accounts for the

[28]Results 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.

fact that poorer Americans and minorities are, overall, substantially more focused on efficiency than the rest of the population.

7 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 in formulating, and understanding support for, a range of social and redistributive 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 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 Stata 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.

Becker, G. S. (1962): "Irrational Behavior and Economic Theory," *The Journal of Political Economy*, 70(1), 1–13.

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.

Bronars, S. G. (1987): "The Power of Nonparametric Tests of Preference Maximization," *Econometrica*, 55(3), 693–698.

Camerer, C. (2003): *Behavioral Game Theory*. Princeton University Press, Princeton, NJ.

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.

Choi, S., S. Kariv, W. Müller, and D. Silverman (forthcoming): "Who Is (More) Rational?," *American Economic Review*.

Croson, R., and U. Gneezy (2009): "Gender Differences in Preferences," *Journal of Economic Literature*, 47(2), 1365–1384.

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*.

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.

Kuziemko, I. (in progress): "Redistributive preferences in an age of rising inequality," upublished 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.

Rawls, J. (1971): *A Theory of Justice*. Harvard University Press, Cambrudge, MA.

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: Subjects' Progression through the Experiment

Logged in to experiment	1.00
Started incentivized decision problems	0.89
Completed entire experiment	0.85

Includes data on 1,172 ALP respondents who logged in to the experiment.

Table 3: Summary Statistics on Experimental Subjects

Variable:	Mean	S.D.	Median	Min.	Max.	N
Female	0.58	0.49	.	.	.	1002
Age	49.12	15.02	50	19	91	1002
Completed high school	0.91	0.29	.	.	.	1002
Some college	0.29	0.45	.	.	.	1002
Completed college	0.31	0.46	.	.	.	1002
Married	0.61	0.49	.	.	.	1002
Widowed, separated, or divorced	0.21	0.41	.	.	.	1002
African American	0.11	0.31	.	.	.	1000
Hispanic or Latino	0.18	0.38	.	.	.	1002
Non-Hispanic white	0.68	0.47	.	.	.	1000
Protestant	0.24	0.43	.	.	.	994
Catholic	0.24	0.43	.	.	.	994
Not religious	0.23	0.42	.	.	.	994
Employed	0.56	0.50	.	.	.	1002
Unemployed	0.10	0.30	.	.	.	1002
Retired	0.17	0.38	.	.	.	1002
Disabled	0.08	0.27	.	.	.	1002
Homemaker	0.06	0.24	.	.	.	1002
HH income (thousands of USD)	54.98	46.43	45	2.5	206.25	998
Lives in northeast (census region I)	0.18	0.38	.	.	.	1002
Lives in midwest (census region II)	0.20	0.40	.	.	.	1002
Lives in south (census region III)	0.35	0.48	.	.	.	1002
Lives in west (census region IV)	0.27	0.44	.	.	.	1002

Data on race, household income, and religious affiliation is missing for (respectively) 2, 5, and 8 respondents.

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

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(20)
Female	-0.023*	.	.	.	.	.	.	.	-0.015	-0.021
	(0.013)								(0.014)	(0.014)
Youngest quartile (age 37 or less)	.	-0.008	.	.	.	.	.	.	-0.003	-0.004
		(0.015)							(0.016)	(0.017)
Oldest quartile (over 60)	.	0.035**	.	.	.	.	.	.	0.026	0.025
		(0.016)							(0.018)	(0.018)
Did not complete high school	.	.	-0.054***	.	.	.	.	.	-0.046**	-0.039**
			(0.019)						(0.019)	(0.02)
Completed college	.	.	0.029*	.	.	.	.	.	0.009	0.007
			(0.016)						(0.016)	(0.016)
African American	.	.	.	-0.08***	.	.	.	.	-0.063***	-0.066***
				(0.018)					(0.019)	(0.02)
Hispanic/Latino	.	.	.	-0.047***	.	.	.	.	-0.018	-0.017
				(0.015)					(0.017)	(0.019)
Lowest income quartile	.	.	.	.	-0.033**	.	.	.	-0.0004	-0.002
					(0.014)				(0.017)	(0.017)
Highest income quartile	.	.	.	.	0.012	.	.	.	-0.002	-0.002
					(0.017)				(0.018)	(0.018)
Employed	.	.	.	.	.	0.009	.	.	0.003	0.005
						(0.018)			(0.017)	(0.017)
Unemployed	.	.	.	.	.	-0.048**	.	.	-0.026	-0.03
						(0.023)			(0.023)	(0.023)
Married	.	.	.	.	.	.	0.037**	.	0.002	-0.004
							(0.017)		(0.019)	(0.019)
Widowed, separated, or divorced	.	.	.	.	.	.	0.012	.	-0.016	-0.011
							(0.02)		(0.021)	(0.022)
Catholic	.	.	.	.	.	.	.	-0.022	-0.029	-0.038*
								(0.017)	(0.018)	(0.019)
Protestant	.	.	.	.	.	.	.	0.031*	0.006	-0.003
								(0.018)	(0.018)	(0.019)
No religious preference	.	.	.	.	.	.	.	-0.014	-0.018	-0.016
								(0.018)	(0.018)	(0.018)
Constant	0.69***	0.67***	0.664***	0.694***	0.684***	0.672***	0.651***	0.677***	0.704***	0.714***
	(0.01)	(0.009)	(0.011)	(0.008)	(0.009)	(0.016)	(0.015)	(0.012)	(0.029)	(0.03)
State of Residence FEs	No	No	No	No	No	No	No	No	No	Yes
Observations	1002	1002	1002	1002	1002	1002	1002	1002	1002	1002
R^2	0.003	0.007	0.016	0.022	0.008	0.009	0.006	0.01	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 5: OLS Regressions of Estimated $\hat{\rho}_n$ on Subject Characteristics

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Female	-0.773**	.	.	.	.	.	.	.	-0.94**	-1.062***
	(0.377)								(0.396)	(0.402)
Youngest quartile (age 37 or less)	.	1.512***	.	.	.	.	.	.	1.418***	1.457***
		(0.396)							(0.414)	(0.43)
Oldest quartile (over 60)	.	-0.413	.	.	.	.	.	.	0.017	-0.096
		(0.516)							(0.599)	(0.603)
Did not complete high school	.	.	0.562	.	.	.	.	.	0.057	0.417
			(0.711)						(0.673)	(0.684)
Completed college	.	.	-0.082	.	.	.	.	.	-0.096	-0.375
			(0.5)						(0.469)	(0.485)
African American	.	.	.	1.042*	.	.	.	.	0.747	0.412
				(0.557)					(0.672)	(0.722)
Hispanic/Latino	.	.	.	0.958**	.	.	.	.	0.111	0.171
				(0.452)					(0.551)	(0.624)
Lowest income quartile	.	.	.	.	0.823*	.	.	.	1.137**	1.041**
					(0.423)				(0.512)	(0.525)
Highest income quartile	.	.	.	.	-0.441	.	.	.	-0.622	-0.722
					(0.517)				(0.533)	(0.532)
Employed	.	.	.	.	.	0.694	.	.	0.918*	0.651
						(0.556)			(0.543)	(0.53)
Unemployed	.	.	.	.	.	1.018	.	.	0.372	0.003
						(0.741)			(0.725)	(0.731)
Married	.	.	.	.	.	.	-1.001**	.	-0.013	-0.006
							(0.451)		(0.503)	(0.522)
Widowed, separated, or divorced	.	.	.	.	.	.	-1.471**	.	-0.644	-0.33
							(0.59)		(0.626)	(0.658)
Catholic	.	.	.	.	.	.	.	0.752	0.856	0.603
								(0.509)	(0.58)	(0.583)
Protestant	.	.	.	.	.	.	.	-0.247	0.283	0.369
								(0.551)	(0.576)	(0.595)
No religious preference	.	.	.	.	.	.	.	0.067	-0.132	-0.217
								(0.551)	(0.542)	(0.555)
Constant	-2.188***	-2.939***	-2.786***	-2.908***	-2.791***	-3.062***	-1.719***	-2.796***	-3.408***	-3.013***
	(0.27)	(0.285)	(0.345)	(0.233)	(0.284)	(0.501)	(0.378)	(0.371)	(0.841)	(0.834)
State of Residence FEs	No	No	No	No	No	No	No	No	No	Yes
Observations	1002	1002	1002	1002	1002	1002	1002	1002	1002	1002
R^2	0.004	0.015	0.002	0.009	0.007	0.006	0.006	0.005	0.044	0.107

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: Additional Regressions of Estimated $\hat{\rho}_n$ on Subject Characteristics

Dependent Variable:	Estimated $\hat{\rho}_n$		Estimated $\hat{\rho}_n$		Decile of $\hat{\rho}_n$		ρ_{high}	
Specification:	OLS Regression		Median Regression		OLS Regression		OLS Regression	
Subjects included:	All	Non-Selfish	All	Non-Selfish	All	Non-Selfish	All	Non-Selfish
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Female	-0.94**	-0.876**	-0.137**	-0.117*	-0.497***	-0.381**	-0.066**	-0.048
	(0.396)	(0.415)	(0.06)	(0.063)	(0.191)	(0.191)	(0.032)	(0.033)
Youngest quartile (age 37 or less)	1.418***	1.387***	0.163**	0.149*	0.675***	0.594***	0.088**	0.081**
	(0.414)	(0.428)	(0.074)	(0.079)	(0.226)	(0.224)	(0.04)	(0.041)
Oldest quartile (over 60)	0.017	0.087	-0.095	-0.085	-0.237	-0.199	-0.06	-0.055
	(0.599)	(0.621)	(0.081)	(0.085)	(0.262)	(0.261)	(0.044)	(0.044)
Did not complete high school	0.057	0.119	0.123	0.133	0.38	0.498	0.101*	0.118**
	(0.673)	(0.678)	(0.107)	(0.11)	(0.326)	(0.326)	(0.058)	(0.059)
Completed college	-0.096	-0.119	0.144**	0.147**	0.295	0.278	0.046	0.047
	(0.469)	(0.49)	(0.07)	(0.074)	(0.22)	(0.218)	(0.037)	(0.038)
African American	0.747	0.898	0.313***	0.339***	0.657**	0.827***	0.151***	0.176***
	(0.672)	(0.683)	(0.1)	(0.104)	(0.32)	(0.321)	(0.055)	(0.056)
Hispanic/Latino	0.111	0.204	0.042	0.062	-0.086	-0.01	0.019	0.031
	(0.551)	(0.566)	(0.084)	(0.088)	(0.266)	(0.263)	(0.046)	(0.046)
Lowest income quartile	1.137**	1.220**	0.259***	0.239***	0.484**	0.525**	0.078*	0.084**
	(0.512)	(0.529)	(0.077)	(0.08)	(0.238)	(0.238)	(0.042)	(0.042)
Highest income quartile	-0.622	-0.658	0.005	-0.039	-0.013	-0.003	0.031	0.036
	(0.533)	(0.564)	(0.08)	(0.085)	(0.253)	(0.253)	(0.043)	(0.044)
Employed	0.918*	0.93	0.128*	0.104	0.025	0.006	-0.009	-0.015
	(0.543)	(0.567)	(0.074)	(0.078)	(0.244)	(0.242)	(0.04)	(0.04)
Unemployed	0.372	0.375	0.044	0.04	-0.015	-0.002	0.003	0.005
	(0.725)	(0.747)	(0.109)	(0.114)	(0.332)	(0.329)	(0.058)	(0.059)
Married	-0.013	0.023	0.011	0.006	-0.173	-0.112	-0.028	-0.021
	(0.503)	(0.527)	(0.086)	(0.091)	(0.267)	(0.262)	(0.047)	(0.048)
Widowed, separated, or divorced	-0.644	-0.613	-0.111	-0.106	-0.335	-0.251	-0.039	-0.027
	(0.626)	(0.641)	(0.099)	(0.104)	(0.311)	(0.307)	(0.055)	(0.055)
Catholic	0.856	0.996*	0.107	0.153*	0.46*	0.655**	0.056	0.082*
	(0.58)	(0.601)	(0.084)	(0.088)	(0.274)	(0.273)	(0.046)	(0.046)
Protestant	0.283	0.349	-0.126	-0.139	-0.166	-0.076	-0.012	-0.006
	(0.576)	(0.604)	(0.082)	(0.087)	(0.259)	(0.259)	(0.044)	(0.045)
No religious preference	-0.132	-0.034	-0.135*	-0.118	-0.119	-0.02	-0.012	-0.00009
	(0.542)	(0.565)	(0.081)	(0.085)	(0.259)	(0.258)	(0.044)	(0.045)
Constant	-3.408***	-3.749***	-0.382***	-0.409***	5.455***	5.045***	0.394***	0.336***
	(0.841)	(0.889)	(0.127)	(0.135)	(0.408)	(0.4)	(0.069)	(0.07)
Observations	1002	957	1002	957	1002	957	1002	957

Standard errors in parentheses (robust standard errors in Columns 1 and 2). All regressions include controls for respondents who are missing data on race (2), household income (5), or religious affiliation (8). ρ_{high} is an indicator for being efficiency-focused in the sense of having $\hat{\rho}_n$ greater than or equal to 0. Non-Selfish subjects are those who allocate themselves less than 99 percent of the tokens, on average.

Table 7: OLS Regressions of Likelihood of Voting for Obama in 2012

	— Without State FEs —			— With State FEs —		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: All Subjects						
$\hat{\rho}_n$	-0.003	.	.	-0.005*	.	.
	(0.003)			(0.003)		
Decile of $\hat{\rho}_n$	.	-0.009	.	.	-0.013**	.
		(0.006)			(0.006)	
ρ_{high} (i.e. $\hat{\rho}_n \geq 0$)	.	.	-0.045	.	.	-0.068**
			(0.033)			(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
Panel B: Non-Selfish Subjects						
$\hat{\rho}_n$	-0.004	.	.	-0.006*	.	.
	(0.003)			(0.003)		
Decile of $\hat{\rho}_n$	.	-0.012**	.	.	-0.016**	.
		(0.006)			(0.006)	
ρ_{high} (i.e. $\hat{\rho}_n \geq 0$)	.	.	-0.057*	.	.	-0.077**
			(0.034)			(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 8: OLS Regressions of Likelihood of Being a Democrat

	— Without State FEs —			— With State FEs —		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: All Subjects						
$\hat{\rho}_n$	-0.002	.	.	-0.005	.	.
	(0.003)			(0.003)		
Decile of $\hat{\rho}_n$	.	-0.014*	.	.	-0.02***	.
		(0.007)			(0.007)	
ρ_{high} (i.e. $\hat{\rho}_n \geq 0$)	.	.	-0.075*	.	.	-0.104**
			(0.04)			(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
Panel B: Non-Selfish Subjects						
$\hat{\rho}_n$	-0.003	.	.	-0.005	.	.
	(0.003)			(0.003)		
Decile of $\hat{\rho}_n$	.	-0.016**	.	.	-0.023***	.
		(0.007)			(0.008)	
ρ_{high} (i.e. $\hat{\rho}_n \geq 0$)	.	.	-0.087**	.	.	-0.112**
			(0.041)			(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).

Figure 1: Prototypical Fair-minded Distributional Preferences

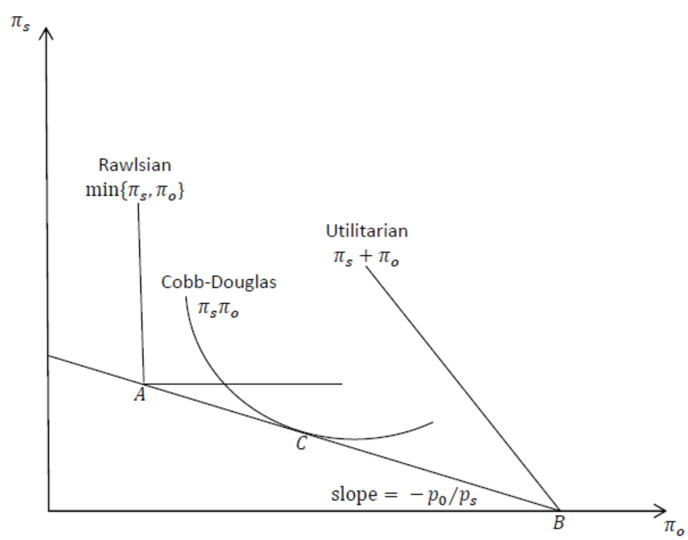

Figure 2: Average Fraction of Tokens Allocated to *Self* $(\pi_s/(\pi_s+\pi_o)$ by Sub-Group

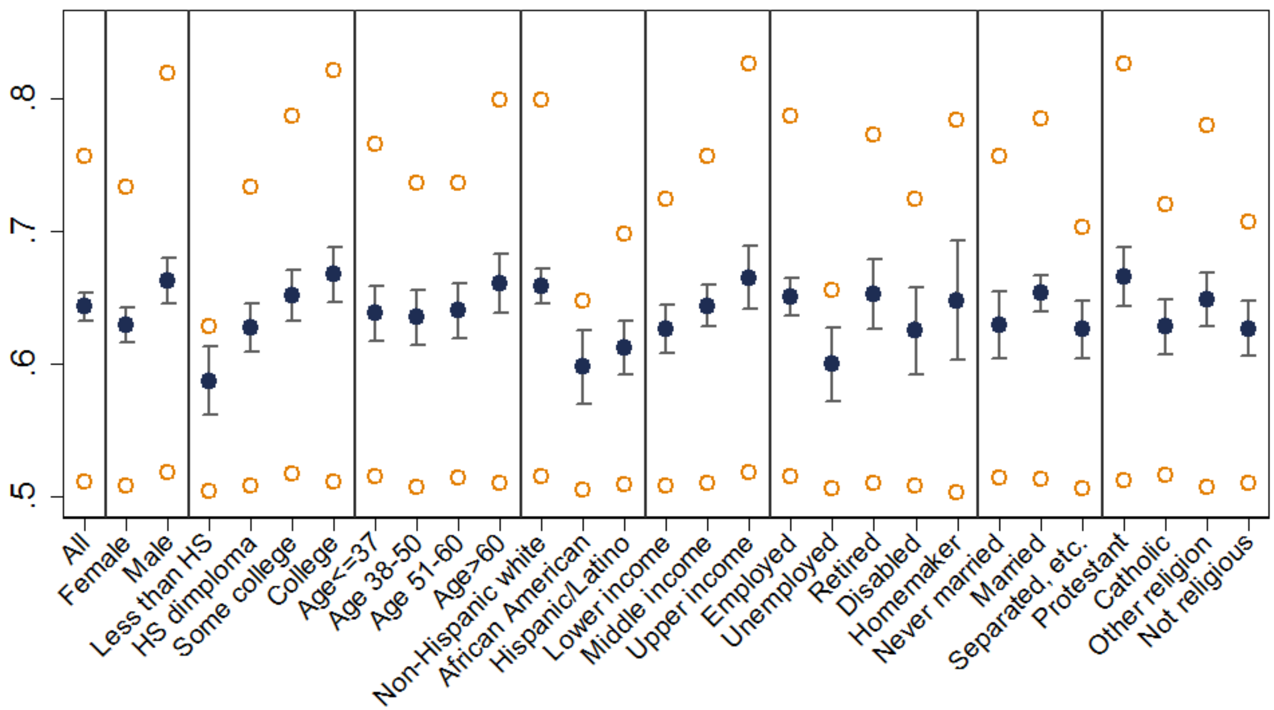

Dots indicate mean values. Circles indicate 25th and 75th percentiles. Bars indicate 95 percent confidence intervals for means.

Figure 3: Proportion of Efficiency-Focused Subjects, by Sub-Group

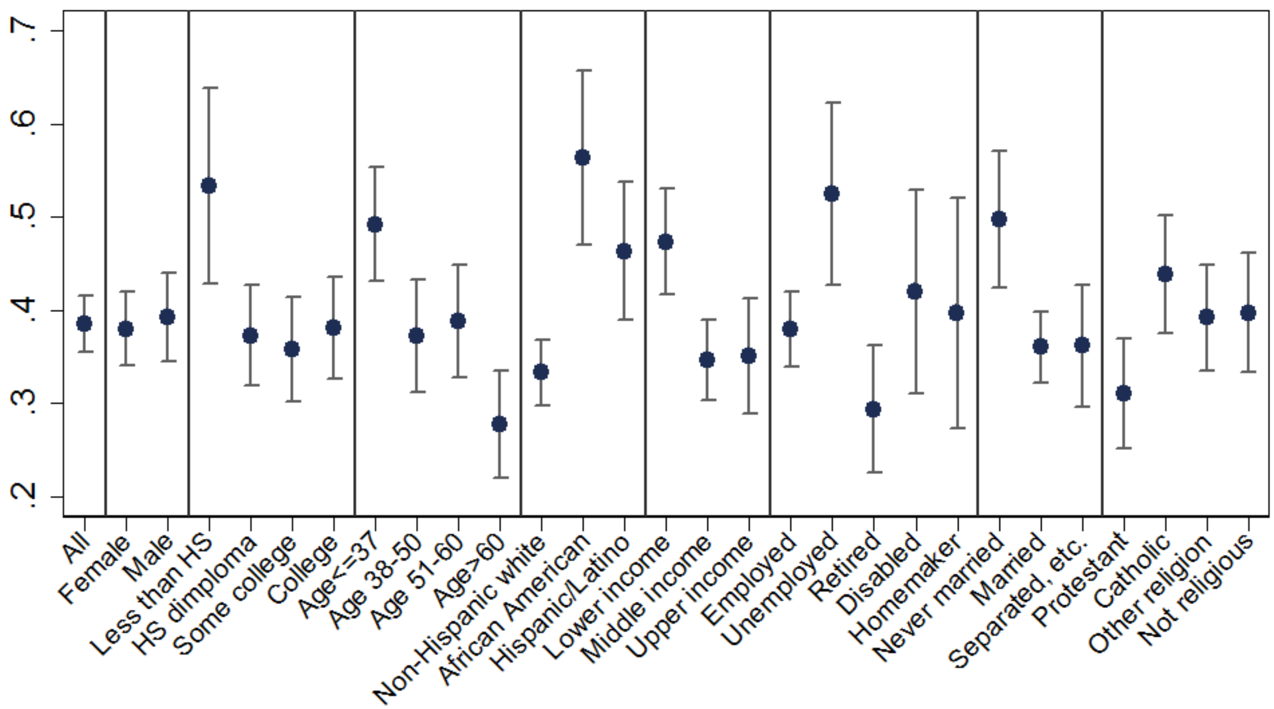

Dots indicate mean values. Bars indicate 95 percent confidence intervals.

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

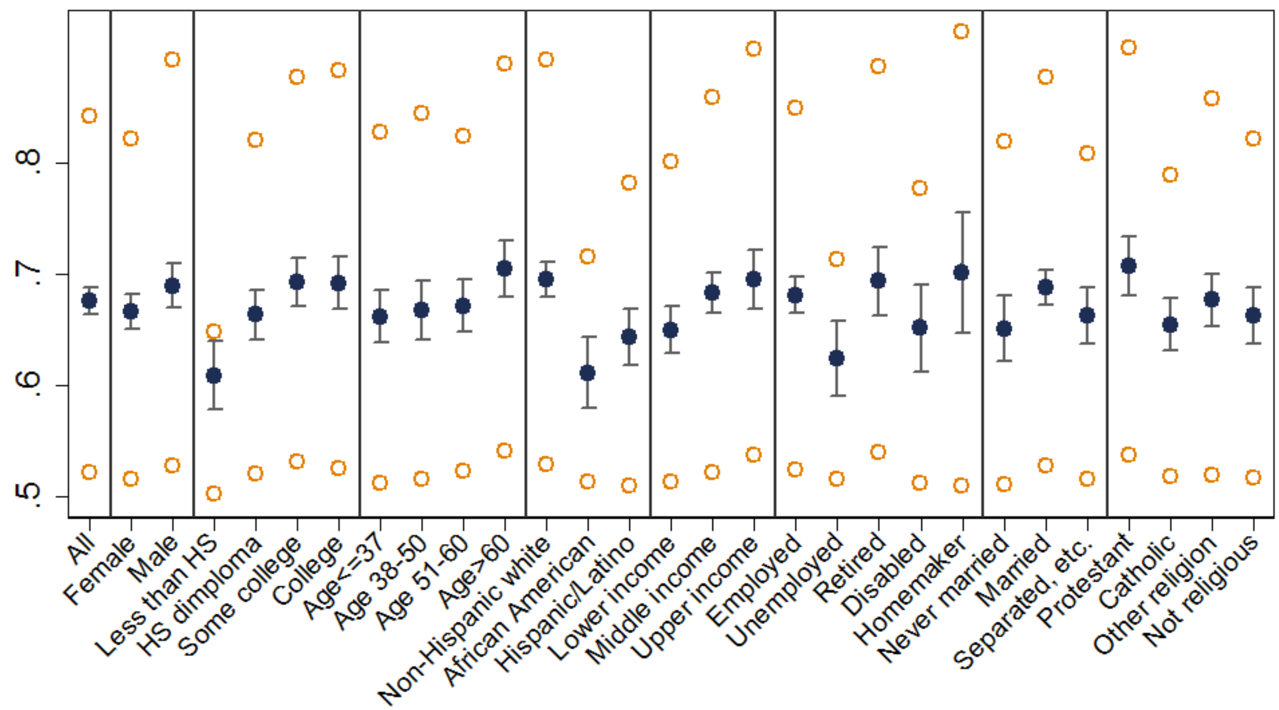

Dots indicate mean values. Circles indicate 25^{th} and 75^{th} percentiles. Bars indicate 95 percent confidence intervals for means.

Figure 5: Estimated Mean $\hat{\rho}_n$ Parameters, by Sub-Group

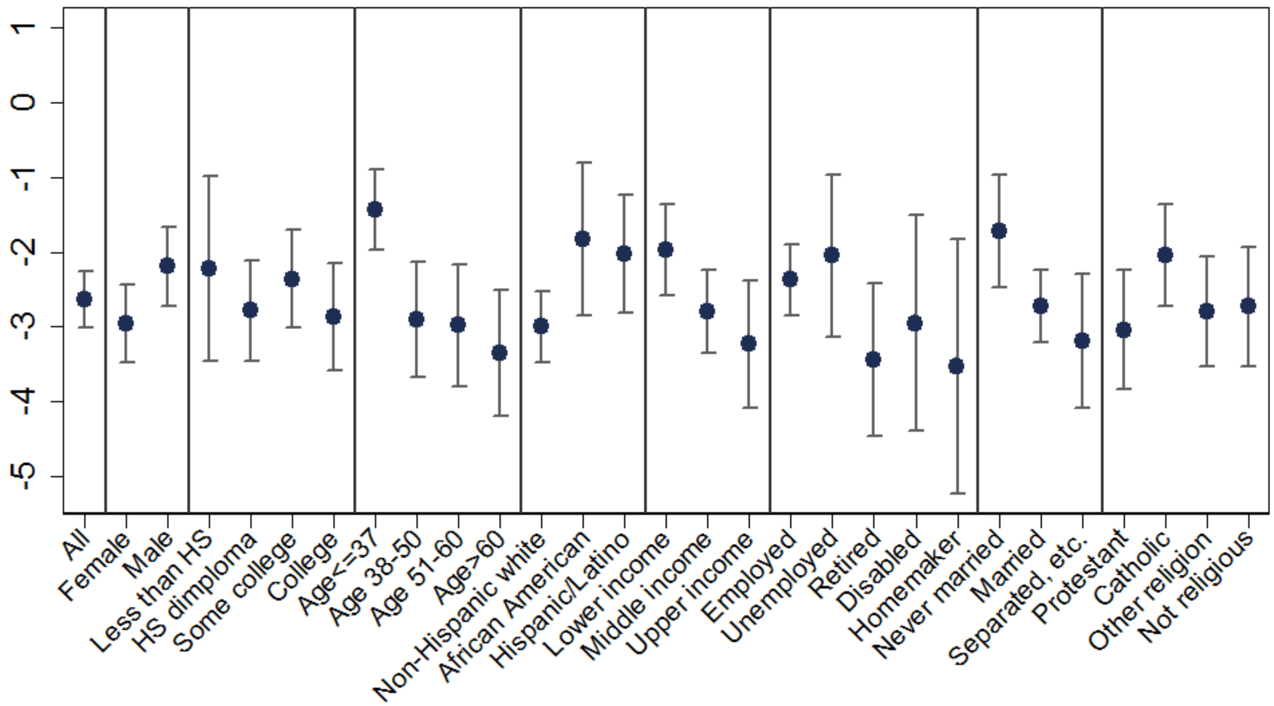

Dots indicate mean values. Bars indicate 95 percent confidence intervals for means.

Figure 6: Estimated Median $\hat{\rho}_n$ Parameters, by Sub-Group

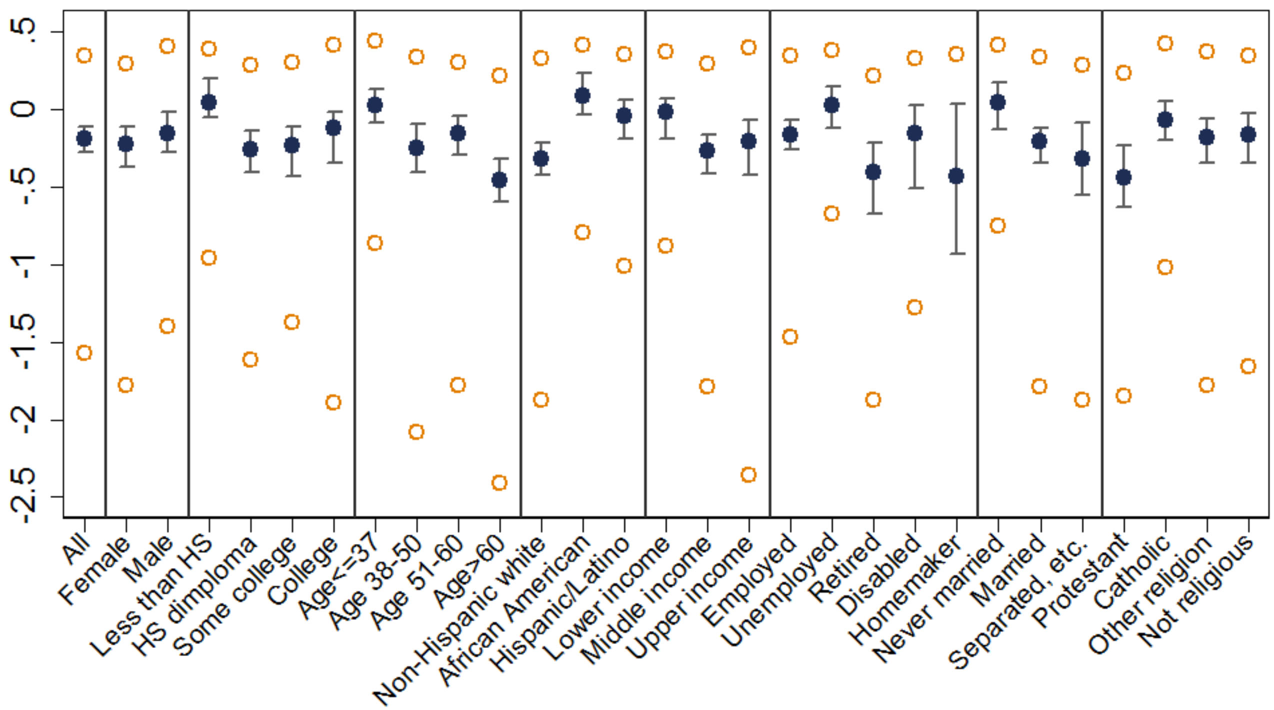

Dots indicate median values. Circles indicate 25^{th} and 75^{th} percentiles. Bars indicate 95 percent confidence intervals for medians.