

# The Distributional Preferences of Americans, 2013-2016\*

Raymond Fisman, Pamela Jakiela, Shachar Kariv and Silvia Vannutelli<sup>†</sup>

Oct 22, 2020

## Abstract

We study the distributional preferences of Americans during 2013-2016, a period of social and economic upheaval. We decompose preferences into two qualitatively different tradeoffs – fairness versus self-interest, and equality versus efficiency – and measure both at the individual level in a large and diverse sample. Although Americans are heterogeneous in terms of both fair-mindedness and equality-efficiency orientation, we find that the individual-level preferences in 2013 are highly predictive of those in 2016. Subjects that experienced an increase in household income became more self-interested, and those who voted for Democratic presidential candidates in both 2012 and 2016 became more equality-oriented.

**JEL Classification Numbers:** C79, C91, D63.

**Keywords:** distributional preferences, social preferences, fairness, impartiality, equality, efficiency, rationality, revealed preference, redistribution, political decisions, voting, household income, American Life Panel (ALP), experiment.

---

\*We are grateful to Ernst Fehr and Daniel Markovits for helpful discussions and for suggestions from the audience at the Ben-Porath Annual Lecture at the Hebrew University of Jerusalem and a number of seminar participants. We thank the American Life Panel (ALP) team at the RAND Corporation for technical and administrative support. We acknowledge financial support from the Center for Equitable Growth (CEG) at the University of California, Berkeley. Any findings, opinions and conclusions expressed in this material are those of the authors and do not necessarily reflect the views of the institutions or any funding agency.

<sup>†</sup>Fisman: Boston University (rfisman@bu.edu); Jakiela: Center for Global Development (pjakiela@cgdev.org); Kariv: University of California, Berkeley (kariv@berkeley.edu); Vannutelli: Boston University (svann@bu.edu).

# 1 Introduction

In many social and economic situations, problems of cooperation and conflict arise in connection with the distribution of resources under scarcity. These problems emerge not only because people promote their competing private interests but also because fair-minded people who are concerned with the interests of others will often disagree about what fairness requires, especially regarding equality versus efficiency.<sup>1</sup> As a result, distributional preferences – the tradeoff between fair-mindedness and self-interest, and the tradeoff between equality and efficiency – shape individual opinions on a range of issues related to the redistribution of income, including social security, unemployment benefits, and government-sponsored healthcare. The theoretical and empirical analysis of these preferences therefore has broad-reaching implications not just for economic policy but also for a host of applications beyond economics. Experimental economics has been very fruitful in both establishing the empirical reliability of distributional preferences – testing for the consistency of such preferences and identifying their underlying structure – and directing theoretical attention to such preferences.<sup>2</sup> Our prior work finds significant heterogeneity across individuals’ distributional preferences, correlates these preferences with sociodemographic categories and political ideology (Fisman et al., 2017), and further relates these preferences with, for example, macroeconomic conditions (Fisman et al., 2015a) and behaviors outside the laboratory (Fisman et al. 2015b; Li et al. 2017). But the extent to which distributional preferences are fixed or malleable is a much less well-studied question, and one we take up here.

In this paper we study the distributional preferences of a large and diverse sample of Americans during the years 2013 – 2016, a period during which economic inequality has continued to increase in the U.S. and elsewhere (Zucman, 2019) and the U.S. has experienced particularly strong economic, cultural, and political upheaval. This makes it an opportune mo-

---

<sup>1</sup>In a classic series of writings, Harsanyi (1955) and Rawls (1971) construct theories of distributional preferences based on the choices that people would or should make for society in what Rawls terms the “original position.” As Harsanyi (1955) and Harsanyi (1978) points out, he and Rawls agree on fair-mindedness (satisfying impartiality and impersonality) but come to quite different conclusions how distributional preferences should tradeoff equality and efficiency. The work of Harsanyi and Rawls, and of the many others who have followed them, has had broad-reaching influence across disciplines.

<sup>2</sup>We will not attempt to review the large and growing theoretical, empirical and experimental literatures on distributional preferences. Camerer (1995), Camerer (2003) and Cooper and Kagel (2016) provide a comprehensive discussion of experimental and theoretical work in economics focusing on dictator, ultimatum and trust games.

ment to study whether and how individuals' distributional preferences have changed. To our knowledge this is the first paper to study the intertemporal stability of preferences – distributional or otherwise – across several years in a large-scale experiment with a broadly representative subject pool.

We emphasize that both fair-mindedness as well as equality-efficiency orientation are potentially implicated by events during the period we study:

- **Fairness vs. self-interest** The 2016 election unsettled conventional associations between partisan politics and economic policy. A populist tide suffused both parties and overcame the Republican establishment to produce, in Donald Trump, a presidential candidate who rejected traditional Republican efficiency-minded policy positions, for example on trade and top tax rates. Trump's populism also had an economic face that both responded to and addressed the material strains and dislocations that globalization and technological change have imposed on middle class workers. Indeed, renewed economic self-interest by a disenchanted middle class constitutes one of the 2016 election's leading themes.
- **Equality vs. efficiency** The major parties in the U.S. have, in recent decades, competed expressly with respect to equality-efficiency tradeoffs. The Republican Party, by championing lower taxes and emphasizing business, has embraced efficiency, while the Democratic Party, by defending higher taxes and emphasizing the social welfare state, has embraced equality. The abstract and general programs have played out concretely, in specific policy disputes during the period we study, including in the debate about whether or not to decrease the top marginal tax rate and in the battle over the taxes and subsidies contained in the Affordable Care Act. As partisanship hardened in the years leading up to the 2016 election (Doherty et al., 2016), it is possible that equality-efficiency orientations have similarly diverged.<sup>3</sup>

To study changes in distributional preferences, we conduct identical experiments three years apart – in 2013 and 2016 – using the American Life Panel (ALP), a longitudinal survey administered online to more than 5,000

---

<sup>3</sup>As Kuziemko et al. (2015) concluded from survey-based evidence, at least some of the partisan gap in attitudes toward taxation and redistribution may result from attitudes towards government in general, rather than underlying distributional preferences. Our subjects make direct tradeoffs between their own payoff and the payoff an individual drawn from the general population, so their equality-efficiency orientation is not confounded by their beliefs about the role of government and its effectiveness in redistributing income.

Americans. 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 American adults, and track their changes over time. 687 ALP respondents who completed our experiment in 2013 and 2016 constitute our subject pool.

In our experiment, we presented subjects with a sequence of decision problems – each depicted as a choice from a two-dimensional budget line represented graphically on a computer screen. Every subject completed the experiment twice, once in 2013 and again in 2016. The subjects used the mouse to select an allocation from the budget line by pointing and clicking. Each allocation from the budget line represents an allocation of monetary payoffs to *self* (the subject) herself and an anonymous *other* (an ALP respondent not sampled for the experiment). The budget line is given by  $p_s\pi_s + p_o\pi_o$  where  $\pi_s$  and  $\pi_o$  correspond to the payoffs of *self* and *other*, and  $p_s/p_o$  is the relative price of redistribution. Because choices are from standard budget lines, we are able to use classical revealed preference analysis to assess whether subjects’ behavior is consistent with rationality, and classical demand analysis to recover information about the underlying preferences, specifically to decompose distributional preferences into fair-mindedness and equality-efficiency orientation.

We begin our analysis of the experimental data by using revealed preference theory (Afriat 1967; Varian 1982; Varian 1983) to determine whether the observed choices are consistent with *self* maximizing a utility function  $u_s(\pi_s, \pi_o)$  that captures distributional preferences. We find that, for most subjects, the violations are sufficiently minor that we can ignore them for the purposes of constructing appropriate utility functions and recovering preferences. We then use the experimental data to estimate the underlying distributional preferences at the level of the individual subject  $n$  in 2013 and 2016.

Following Andreoni and Miller (2002) and Fisman et al. (2007), we estimate these preferences using constant CES utility function for each subject:

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

where  $\alpha$  is the weight on  $\pi_s$  versus  $\pi_o$  (fair-mindedness) and  $\rho$  is the weight on reducing the difference between  $\pi_s$  versus  $\pi_o$  (equality) versus increasing the sum of  $\pi_s$  and  $\pi_o$  (efficiency).  $\alpha = 1/2$  indicates pure fair-mindedness (im-

partial treatment of *self* vis-a-vis *other*) whereas  $\alpha = 1$  indicates pure self-interest; any  $0 < \rho \leq 1$  indicates distributional preference weighted towards efficiency (increasing total payoffs) because the expenditure on the tokens kept by *self*  $p_s \pi_s$  decreases when the relative price  $p_s/p_o$  increases, whereas any  $\rho < 0$  indicates distributional preference weighted towards equality (reducing differences in payoffs) because  $p_s \pi_s$  increases when  $p_s/p_o$  increases.

We use the individual-level CES estimates of fair-mindedness  $\hat{\alpha}_n$  and equality-efficiency orientation  $\hat{\rho}_n$  to compare the distributional preferences of ALP participants in 2013 and 2016. Table 1 provides a summary of the individual-level estimates  $\hat{\alpha}_n$  and  $\hat{\rho}_n$  using a 5% significance level. To provide a snapshot of our main results, we first classify subjects as purely self-interested ( $\hat{\alpha}_n = 1$ ) or fair-minded ( $\hat{\alpha}_n < 1$ ), and then classify the fair-minded subjects as equality-focused ( $\hat{\rho}_n < 0$ ), intermediate ( $\hat{\rho}_n = 0$ ) or efficiency-focused ( $\hat{\rho}_n > 0$ ). Note that the equality-efficiency orientation of purely self-interested subjects, who always allocate almost nothing to *other*, cannot be cleanly identified (Section 5 provides more detail on this categorization). The right column and bottom row of Table 1 report the preference classifications in 2013 and 2016, respectively. Subjects whose classification remained unchanged between 2013 and 2016 are captured along the top-left to bottom-right diagonal.

[Table 1 here]

At the population-level, Table 1 shows a high degree of similarity in the classifications of distributional preferences in 2013 and 2016 – a period of rising economic inequality and deep frustration and resentment among ordinary Americans, that might very plausibly lead to substantial changes in distributional preferences. Specifically, the fractions of purely self-interested subjects ( $\hat{\alpha}_n = 1$ ) and equality-focused subjects ( $\hat{\rho}_n < 0$ ) remain nearly unchanged (14.0% vs. 14.9% and 23.2% vs. 22.3%, respectively) whereas the fraction of efficiency-focused subjects ( $\hat{\rho}_n > 0$ ) increased (from 25.2% to 31.9%) and the fraction of subjects with intermediate equality-efficiency orientation ( $\hat{\rho}_n = 0$ ) correspondingly decreased (from 37.6% to 31.0%). One might expect that the continuing increase in economic inequality in the U.S. in the period 2013-2016 would increase the willingness to sacrifice economic efficiency to combat income inequality, but our results show that, if anything, the Americans in our sample put a higher emphasis on efficiency relative to equality in 2016 versus 2013.

At the individual-level, of our 687 subjects, 45.3% have the same preference classification in 2013 and 2016. We observe shifts in fair-mindedness

for 20.5% of subjects – 10.7% from pure self-interest ( $\hat{\alpha}_n = 1$ ) to fair-minded ( $\hat{\alpha}_n < 1$ ) and 9.8% in the other direction – and shifts in equality-efficiency orientation for 34.3% of subjects. The latter set of shifts tend to be to or from the intermediate category ( $\hat{\rho}_n = 0$ ) rather than extreme shifts – 2.5% of subjects shift from efficiency-focused ( $\hat{\rho}_n > 0$ ) to equality-focused ( $\hat{\rho}_n < 0$ ) and 3.3% in the other direction, a total of only 5.8% of subjects. (We obtain near-identical results if we use other conventional significance levels or classifications based on point estimates instead of statistical classification.)

The ALP collects extensive demographic and economic information from its members so we can combine our rich individual-level experimental data with the individual demographic and economic information from the survey to study if and how changes in distributional preferences are explained by personal attributes and/or life events. We focus on two main factors that we see as the most plausible sources of shifting distributional preferences during the 2013-2016 time-period – economic circumstances and political preferences.

We begin by looking at changes in household income between 2013 and 2016. We show that large shifts in household income – movements across quartiles of the income distribution – do predict changes in fair-mindedness, in a way that is consistent with the cross-sectional relationship between income and fair-mindedness – higher income is associated with greater self-interest. We find no discernible relationship between income changes and shifts in equality-efficiency orientation. We also evaluate the link between political decisions and distributional preferences and find that subjects who voted for the Democrat candidate in the Presidential election in both 2012 and 2016, reflecting a clear Democratic partisan orientation, become more equality-oriented. Perhaps intriguingly, we find that shifting political allegiances is not associated with a shift in distributional preferences.

The rest of the paper is organized as follows. Section 2 discusses the related literature. Section 3 describes the subject pool and Section 4 describes the experimental design and empirical methods. Section 5 provides the empirical analysis and results, and Section 6 concludes by discussing the results and relating them to the broader literature.

## 2 Related Literature

Experimental research has been very fruitful in establishing the empirical reliability of distributional preferences and directing guiding theoretical work in this area. Seminal work includes Levine (1998), Fehr and Schmidt (1999),

Charness and Rabin (2002), and Andreoni and Miller (2002) among others. Camerer (1995), Camerer (2003) and Cooper and Kagel (2016) provide excellent, though now somewhat dated, surveys of the vast body of experimental and theoretical research. But the typical (incentivized) experiment has elicited (relatively few) decisions from small and homogenous samples, often comprised of university students. Because the samples are small and homogenous, the results are rarely linked to demographic variables, or information about subjects' choices outside the laboratory. As a result, external validity is a concern (Levitt and List, 2007).

On the other hand, (nonincentivized) surveys about attitudes toward inequality and redistribution have drawn on large and heterogenous samples. Among others, Kuziemko et al. (2015) analyze how information on income inequality in the U.S. affects attitudes toward taxation and redistribution, and Alesina et al. (2018) investigate how beliefs about intergenerational mobility affect attitudes in the U.S. and Western European countries. However, these studies rely primarily on hypothetical questions, and it is difficult to derive the parameters of underlying distributional preference from the answers to such survey questions.

We thus see studies much untapped promise in research that combines incentivized experiments to reveal distributional preferences and survey research. To our knowledge, Fisman et al. (2017), Bellemare et al. (2008) and Almås et al. (2020) are the only papers that study distributional preferences by combining laboratory methods and survey research with broadly representative subject pools. Building on Fisman et al. (2017), this paper is first paper to examine the intertemporal stability of distributional preferences across several years in a large-scale experiment with a broadly representative subject pool, using a methodology that is constant across time and, as our results indicate, sufficiently effective at capturing individual preferences to assess the existence of stability across several years.<sup>4</sup>

Our work thus contributes most directly to the literature that aims to assess the intertemporal stability of distributional preferences, and also con-

---

<sup>4</sup>Prior work that has measured distributional preferences across time is lacking in one or more of these features. The most immediate antecedent is Bruhin et al. (2019), which also characterizes distributional preferences in a lab setting, and compares responses across multiple points in time, though with a student population (as compared to a diverse national sample), and preferences measured several months (as opposed to several years) apart. Chuang and Schechter (2015) analyze social preferences data gathered in experiments conducted with Paraguayan villagers across a number of years, though the experiments they conduct vary across time periods, and are based on a relatively small number of discrete choices, making it difficult to interpret their findings of relatively unstable preferences.

tributes to the larger body of work on preference stability more broadly. Numerous papers (in economics and psychology) study the stability of attitudes toward risk and time and numerous other papers (in psychology) study the stability of personality traits. Social psychologists find relatively high intertemporal stability of personality traits (Schuerger and Tait, 1982), at least in the relatively short run.

Finally, in Fisman et al. (2015b), Fisman et al. (2017) and Li et al. (2017) we build on Fisman et al. (2007) to study distributional preferences with different samples. In this paper, we build especially on (Fisman et al., 2017) to study distributional preferences across time with subjects drawn from broad cross-sections of the general population. Our finding in this paper that higher income is associated with greater self-interest is in line with Fisman et al. (2015a) which finds – in a population of university students – that fair-mindedness increases on average with the onset of the Great Recession, which brought with it lower incomes and expected future earnings. Those results were based on a repeated cross-section whereas our findings here are based on repeated observations the same individuals drawn from a diverse sample that we follow across time.

### 3 Subject Pool

Subjects in our experiment are respondents in the ALP, an internet survey conducted by the RAND Corporation which is administered to more than 5,000 Americans.<sup>5</sup> Our sample in 2016 is drawn from a set of 1,002 ALP respondents who participated in the identical distributional preference experiment of Fisman et al. (2017) in 2013. Of these, 886 respondents (88.4%) were still part of the ALP in 2016, 726 respondents (81.9%) logged in to the experiment and 687 respondents (94.6%) completed the entire experiment. Thus, the fraction of dropouts who logged in and quit the experiment is very low). The 687 subjects who completed the experiment constitute our subject pool; they represent a diverse set of Americans, spanning a broad range in terms of age, educational attainment, household income, employment status, ethnicity, and place of residence.

Table 2 below provides summary statistics of individual-level characteristics. In column (1), we present the data for the sample of 1,002 ALP respondents who completed the experiment in 2013 (Fisman et al., 2017). In columns (2) and (3), we present the data for the subsample of 687 re-

---

<sup>5</sup>The ALP website provides information on panel composition, demographics, attrition and response rates, sampling weights, and a comparison with other data sources.



spondents who also completed the experiment in 2016 based on the 2013 and 2016 ALP questionnaire. Columns (4) and (5) compare our experimental subjects to the entire ALP sample in 2013 and 2016, and columns (6) and (7) compare them to the American Community Survey (ACS) conducted in 2012 and 2016.

*[Table 2 here]*

Focusing on the 2016 questionnaire, subjects range in age from 22 to 94. 79% of the subjects are Caucasian, 16% are Hispanic or Latino, and 10% are African American. 92% of our subjects have a high school diploma and 34% hold college degrees. 58% of subjects are currently employed, 22% are retired, 8% are disabled, 5% are unemployed, 5% are homemakers, and the remainder are on medical leave or otherwise temporarily absent from the workforce. Noticeably, the fraction of subjects unemployed decreased from 10% in 2013 to 5% in 2016; the ACS also indicates a decline from 6% to 4%. Overall, the data in Table 2 show that our subject pool closely matches the general population in terms of age, place of residence, education, and race, and contains often under-represented groups.

Our sample, however, is less well-off economically relative to the U.S. population – the mean household income is approximately \$70,000 per year and the median is approximately \$55,000 dollars, somewhat lower than the ACS averages. We delve into the income differences between our data and the ACS in greater detail in Table 3, which provides the full income distributions for the 2013 and 2016 ALP samples as well as those from the ACS in 2012 and 2016. Two notable differences between our sample and the ACS emerge – we have a relatively large number of low-income households, which is a direct result of the ALP over-sampling economically vulnerable households, and we have relatively few very high-income households (as relatively few wealthy families are willing to act as survey panel members). But overall, the ALP sample contains relatively large numbers of often under-represented sociodemographic groups.

*[Table 3 here]*

The ALP thus provides a unique opportunity to combine experimental data with sociodemographic and economic variables from the survey. But while the ALP collects extensive demographic and economic attributes from its members, these are almost all time-invariant. In our regression analysis, we focus on two factors – economic circumstances and political preferences

– that vary across individuals, change over time, and are the most plausible determinants of distributional preferences based on our prior work (Fisman et al. 2015a; Fisman et al. 2017):

- To capture substantial shifts in household incomes between 2013 and 2016, we calculate household income quartiles in 2013 and 2016, and use these data to define an indicator variable which takes on values of  $-1, 0, 1$  based on whether the subject’s household income quartile decreased, stayed the same, or increased, respectively.
- To capture shifts in political preferences, we use participants’ voting decisions in the 2012 and 2016 presidential elections and define an indicator variable which takes on values of  $-1, 0, 1$  based on whether the subject shifted to voting Republican, voted the same (or has missing data on voting), or shifted to voting Democrat, respectively.

Of our 687 subjects, the household income quartile stayed the same between 2013 and 2016 for 503 subjects (73.2%), decreased by one quartile for 50 subjects (7.3%) and increased by one quartile for 120 subjects (17.5%). Very few subjects had a more than one quartile change in their household income. For political preferences, we have voting data on the 2012 and 2016 presidential elections for 394 subjects (57.4%). Of those, 151 subjects (38.3%) voted Democrat in both elections, 129 (32.7%) voted Republican in both elections, 25 (6.3%) shifted to voting Republican, and 14 (3.6%) shifted to voting Democrat. The rest voted for other candidates or did not vote in the 2012 and/or 2016 presidential elections.

## 4 Methods

In this section, we first describe the experimental design and procedures, and then proceed to an overview of our empirical framework. Since we are building on expertise we have acquired in previous work, we economize on space and refer the interested reader to Fisman et al. (2007) and Fisman et al. (2017) for more details.

### 4.1 The Experiment

In a split-the-dollar dictator experiment, first introduced by Forsythe et al. (1994), *self* (the subject) divides a dollar between *self* and *other* in any way he wishes such that  $\pi_s + \pi_o = 1$  (without essential loss of generality,

the endowment is normalized to 1). One respect in which this framework is restrictive is that the set of feasible payoff pairs is always the line with a slope of  $-1$ , so that the problem faced by person  $s$  is simply allocating a fixed total income between *self* and *other*. We study a modified dictator game in which *self* must allocate the dollar across  $\pi = (\pi_s, \pi_o)$  at prices  $p = (p_s, p_o)$ , such that  $p_s\pi_s + p_o\pi_o = 1$ . This configuration creates budget sets over  $\pi_s$  and  $\pi_o$  that allow for the thorough testing of observed dictator behavior – specifically the experiments capture subjects’ fair-mindedness (the weight on the payoff of *other*) and equality-efficiency orientation (concerns for reducing differences in payoffs versus increasing total payoffs).

We conducted identical modified dictator game experiments in 2013 and 2016 with a large and diverse sample of ALP respondents. The experiment consisted of 50 independent decision-problems. In each decision problem, *self* was asked to allocate tokens between herself  $\pi_s$  and an anonymous *other*  $\pi_o$ , where *other* was chosen at random from the group of ALP respondents not recruited in the experiment. Each choice involved choosing a point on a line representing the budget constraint over possible token allocations. Choices are made by using the computer mouse to move the pointer on the computer screen and clicking on the desired point.<sup>6,7</sup>

Each decision problem started by having the computer select a budget set randomly from the set of budget sets that intersect with at least one of the axes at 50 or more tokens, but with no intercept exceeding 100 tokens. Each subject thus faced different budget lines in 2013 and 2016. The budget lines selected for each subject in his decision problems were independent of each other and of the budget lines selected for other subjects in their decision problems. At the end of the experiment, the experimental program first randomly selected one decision round from each subject to carry out – *self* received the tokens that she kept in this round  $\pi_s$  and *other* received the tokens that she passed  $\pi_o$ . Payoffs were calculated in terms of tokens and then converted into money. See Fisman et al. (2017) for an extended description of the experimental interface. Full instructions, including the

---

<sup>6</sup>We restricted choices to allocations on the budget constraint so that subjects could not dispose of payoffs. In Fisman et al. (2007), each choice involved choosing a point on a graph representing a budget set over possible allocations. Since most of their subjects had no violations of budget balancedness in all following experiments we restricted choices to allocations on the budget constraint, which simplified the decision problem and made the computer program easier to use.

<sup>7</sup>Because the interface is extremely user-friendly, it is possible to present each subject with many choices in the course of the experiment, yielding a rich individual-level dataset. This makes it possible to analyze behavior at the level of the individual subject, without the need to pool data or assume that subjects are homogenous.

computer program dialog windows are available in the Online Appendix of Fisman et al. (2017).<sup>8,9</sup>

## 4.2 Empirical Framework

**Nonparametric** The most basic question to ask about choice data is whether it is consistent with individual utility maximization. If budget sets are linear (as in our experiment), classical revealed preference theory Afriat (1967) provides a direct test: choices in a finite collection of budget sets are consistent with maximizing a well-behaved (piecewise linear, continuous, increasing, and concave) utility function if and only if they satisfy the Generalized Axiom of Revealed Preference (GARP). Hence, in order to assess whether our data are consistent with utility-maximizing behavior we need to check whether our data satisfy GARP. Because our subjects make choices over a wide range of budget sets, our data provides a strong test of utility maximization.

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. However, individual choices frequently involve at least some errors: subjects may make computational mistakes or execute intended choices incorrectly, or err in other less obvious ways. To account for the possibility of errors, we assess how nearly individual choice behavior complies with GARP by using the Critical Cost Efficiency Index (CCEI) of Afriat (1972), which measures the fraction by which each budget constraint must be shifted inward in order to remove all violations of GARP. By definition, the CCEI is between 0 and 1: values closer to 1 indicate the data are closer to perfect consistency with GARP and hence to perfect consistency with utility maximization.<sup>10</sup>

---

<sup>8</sup>The experimental platform is applicable to other types of individual choice problems. Choi et al. (2007) study risk preferences, and Ahn et al. (2014) extended the work to a settings with ambiguity. Choi et al. (2014) utilize the same experimental technique to study risk preferences in the CentERpanel (a nationally representative Dutch panel).

<sup>9</sup>It is possible that presenting choice problems graphically biases choice behavior in some particular way but there is no evidence that this is the case. For instance, although Fisman et al. (2007) test a much wider range of budget sets than can be tested using the pencil-and-paper questionnaire method of Andreoni and Miller (2002), over all prices, the behavior elicited graphically is consistent with the behavior elicited non-graphically, as well as with the behavior elicited in the split-the-pie dictator games reported in Camerer (2003).

<sup>10</sup>For further details see Fisman et al. (2017). For an overview of revealed preference theory, see Chambers and Echenique (2016) and the papers by Diewert (2012), Varian (2012) and Vermeulen (2012).

An attractive non-parametric approach to test whether the distributional preferences, and hence choice behavior, in 2013 and 2016 are the same (for a given subject) is to combine the data from the two experiments, compute the CCEI for this combined dataset, and compare that number to the CCEI scores in the separate experiments. Intuitively, if a subject's preferences are consistent within an experiment, choices in that experiment should satisfy GARP, while if preferences are the same across both experiments the union of the choices across the two experiments should also satisfy GARP. By definition, the CCEI for the combined data set can be no larger than the minimum of the CCEI scores for the separate data sets (the second dataset can only offer more opportunities for GARP violations), so one measure of the extent to which preferences in 2013 and 2016 coincide is the difference between the CCEI for the combined data set and the minimum of the CCEI for the two separate datasets.

**Parametric** The theorem of Afriat (1967) tells us that if the data satisfy GARP, there exists a well-behaved utility function  $u_s(\pi_s, \pi_o)$  that rationalizes the individual-level data. We additionally will presume that  $u_s(\pi_s, \pi_o)$  is a member of the constant elasticity of substitution (CES) family commonly employed in demand analysis. As noted in the introduction, the CES utility function is given by

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

where  $0 \leq \alpha \leq 1$  is the weight on  $\pi_s$  versus  $\pi_o$  (fairness versus self-interest) and  $\rho \leq 1$  is the weight on reducing the difference  $\pi_s - \pi_o$  (equality) versus increasing the sum  $\pi_s + \pi_o$  (efficiency). Thus,  $\alpha = 1/2$  indicates fair-mindedness (impartiality) whereas  $\alpha = 1$  indicates pure self-interest. The CES approaches a perfect substitutes utility function  $\alpha\pi_s + (1 - \alpha)\pi_o$  as  $\rho \rightarrow 1$  the Leontief form  $\min\{\alpha\pi_s, (1 - \alpha)\pi_o\}$  as  $\rho \rightarrow -\infty$ . As  $\rho \rightarrow 0$ , the indifference curves approach those of a Cobb-Douglas function  $\pi_s^\alpha \pi_o^{1-\alpha}$ .

The CES demand function, expressed as expenditure share, is given by

$$p_s \pi_s = \frac{g}{(p_s/p_o)^r + g} \tag{1}$$

where  $g = [\alpha/(1 - \alpha)]^{1/(1-\rho)}$  and  $r = \rho/(\rho - 1)$ . Note that for any  $\rho > 0$  (resp.  $\rho < 0$ ), a decrease in the relative price  $p_s/p_o$  raises (resp. lowers) the expenditure share on tokens kept by *self*  $p_s \pi_s$ . We generate estimates of  $\hat{g}_n$  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  $\hat{\rho}_n$ .

We emphasize again that our graphical interface enables us to collect many observations per subject, thus allowing us generate parameter estimates for each subject  $n$  separately. Additional details on the estimation may be found in Fisman et al. (2017).

## 5 Results

**Individual Behavior** To convey the wide range of behaviors observed, in Figure 1 we present scatterplots of choices of four subjects. For ease of exposition, we have chosen subjects whose behavior corresponds to one of several prototypical fair-minded distributional preferences in both 2013 and 2016. Collectively they illustrate the striking regularity within subjects and heterogeneity across subjects that is characteristic of much of our data.. For each subject, Figure 1 depicts the relationship between  $\log(p_s/p_o)$  and the tokens kept as a fraction of the sum of the tokens kept and given  $\pi_s/(\pi_s + \pi_o)$ , which examines the sensitivity of choices to changes in the relative price of redistribution. The solid line represents the CES estimation in 2013 and the dotted line represents the analogous CES estimation in 2016.

*[Figure 1 here]*

The first three subjects' distributional preferences in 2013 and 2016 exactly coincide. The first subject (top left) chooses nearly equal allocations  $\pi_s = \pi_o$  in both experiments; this behavior is consistent with maximizing the utility function  $\min\{\pi_s, \pi_o\}$  (Rawlsianism with respect to money). The second subject (top right) allocates all tokens to  $\pi_s$  when  $p_s < p_o$  and to  $\pi_o$  otherwise; this behavior is consistent with maximizing the utility function  $\pi_s + \pi_o$  (utilitarianism with respect to money). The third subject (bottom left) chooses nearly equal expenditures  $p_s\pi_s = p_o\pi_o$ ; this behavior is consistent with maximizing the utility function  $\log \pi_s + \log \pi_o$ . We may contrast these cases with that of the last subject (bottom right), which are consistent with utilitarianism in 2013 but with Rawlsianism in 2016, indicating choices that are consistent *within* each experiment, but differ sharply *across* experiments. These are of course special cases, for which the behaviors are very clear. We also find many intermediate cases, but these are more difficult to discern directly via scatterplots.

**Nonparametric Results** Before calibrating the family of CES utility functions, we first test whether choices can be utility-generated. In our experiments, mean CCEIs across all subjects are 0.872 in 2013 and 0.877

in 2016. We interpret these numbers as confirmation that subjects' choices are generally consistent with utility maximization. As further confirmation, we follow Bronars (1987), which builds on Becker (1962), and compare the behavior of our subjects to that of simulated subjects who randomize uniformly on each budget line. Each of the simulated hypothetical subjects makes 50 choices from randomly generated budget sets in the same way that the human subjects do. Mean CCEIs for a random sample of 25,000 simulated subjects are only 0.600.<sup>11</sup>

The histograms in Figure 2 show the fraction of subjects for whom the lower of the two CCEIs for the separate datasets (black) and the CCEI for the combined dataset (gray) are above different critical values. The horizontal axis shows the critical value of the index and the vertical axis measures the percentage of subjects corresponding to each interval. A measure of the extent to which choice behaviors, and hence distributional preferences, in 2013 and 2016 coincide is the difference between the CCEI of for the combined data set and the minimum of the CCEIs in the separate data sets. Out of 687 subjects, 253 (36.8%) have both CCEI scores in 2013 and 2016 above 0.90 and of those, 149 subjects (21.7%) have both CCEIs above 0.95 (none of simulated random subjects have a CCEI score that high).

In comparison, 138 (20.1%) have a combined CCEI score above 0.90 and of those, 66 subjects (9.6%) have a combined CCEI score above 0.95. For all critical values below 0.95, the difference between the minimum and combined CCEI scores is between 15.7 and 22.7 percentage points. This gap provides a measure of the consistency of distributional preferences across the 2013 and 2016 experiments, accounting for the inconsistencies (violations of the revealed preference conditions) within each experiment. Overall, the patterns in our non-parametric analysis suggest that, while there are some within-experiment inconsistency, subjects exhibit a lot of between-experiment consistency overall.

*[Figure 2 here]*

**Parametric Results** We now return to Table 1 which, recall, provides a population-level summary of the individual-level estimation results. We

---

<sup>11</sup>Alternative measures that have been suggested, including Houtman and Maks (1985) and Varian (1991), yield similar conclusions. The Bronars (1987) test only shows that a significant majority of the subjects did much better than the randomly generated subjects. To test whether utility maximization is the correct model, Fisman et al. (2007) generate a sample of simulated subjects who maximize the CES utility function with an idiosyncratic preference shock that has a logistic distribution and show that a significant majority of the subjects did only slightly worse than an ideal (rational) subject.

categorize a subject as self-interested (an own-payoff maximizer) if we cannot reject that  $\hat{\alpha}_n = 1$  at the 5% level. Otherwise, the subject displays some level of fair-mindedness. We then classify the fair-minded subjects as equality-focused (resp. efficiency-focused) if we reject that  $\hat{\rho}_n > 0$  (resp.  $\hat{\rho}_n < 0$ ) at the 5% level. Subjects that satisfy neither criterion are categorized as having intermediate equality-efficiency orientation. We obtain near-identical regression results if we also categorize the equality-efficiency orientation ( $\hat{\rho}_n$ ) of self-interested subjects ( $\hat{\alpha}_n = 1$ ). These results are presented in Appendix Table A1.

As reported in Table 1, 45.3% of the subjects have the same preference classification in 2013 and 2016. 10.7% shifted from pure self-interest ( $\hat{\alpha}_n = 1$ ) to fair-minded ( $\hat{\alpha}_n < 1$ ) and 9.8% in the other direction, a total of 20.5% of subjects. 34.3% of subjects shifted their equality-efficiency orientation, mostly to or from the intermediate category ( $\hat{\rho}_n = 0$ ). Only 2.5% of subjects shifted from efficiency-focused ( $\hat{\rho}_n > 0$ ) to equality-focused ( $\hat{\rho}_n < 0$ ) and only 3.3% in the other direction. Overall, the fractions of purely self-interested subjects and equality-focused subjects remain nearly unchanged between 2013 and 2016 – 14.0% vs. 14.9% and 23.2% vs. 22.3%, respectively. The fraction of efficiency-focused subjects increased (25.2% vs. 31.9%), with a corresponding decrease in the fraction of subjects with intermediate equality-efficiency orientation (37.6% vs. 31.0%). The regression results presented in Table 4 show that the overall relationship between CES estimates in 2013 and 2016 are statistically significant in a standard regression framework, with  $p < 0.001$  in all specifications.

[Table 4 here]

We can see these relationships in greater detail in the three panels of Figure 3 below. The left and middle panels provide scatterplots of the estimated CES parameters  $\hat{\alpha}_n$  and  $\hat{\rho}_n$ , along with a polynomial regression plot with 95% confidence intervals. For each parameter, we observe a positive monotonic relationship for the entire distribution: a higher parameter value in 2013 is associated with a higher value in 2016. On each axis, we also provide a histogram which shows the distribution of parameter estimates in each year. The distributions are not significantly different using a Kolmogorov–Smirnov test. Since we have estimated a two-parameter utility function, distributional preferences cannot be represented by a single univariate measure. To summarize the distributional preferences of our subjects, we use the expenditure share of the tokens kept by *self*  $0 \leq p_s \pi_s \leq 1$  (prices are normalized by income so that  $p_s \pi_s + p_o \pi_o = 1$ ), which depends on  $\hat{\alpha}_n$  and  $\hat{\rho}_n$



in a non-trivial way according to the CES demand function (1). We observe again a positive and monotonic relationship between the expenditure share  $p_s\pi_s$  in 2013 and 2016. The distributions are again not significantly different using a Kolmogorov–Smirnov test.

*[Figure 3 here]*

**Economic Circumstances, Political Preferences, and Distributional Preferences** We next turn to examining whether the changes in distributional preferences that we do observe are explained by two main factors that might potentially have shifted distributional preferences during this time period – political preferences and economic circumstances. We begin by looking at changes in economic circumstances, for which there is variation between 2013 and 2016 in subjects’ conditions. (As discussed in Section 3, while the ALP collects richly detailed information on household attributes, these are almost all time-invariant.) To capture substantial shifts in household incomes between 2013 and 2016, we calculate household income quartiles in 2013 and 2016, and use these data to define an indicator variable which takes on values of  $-1, 0, 1$  based on whether the subject’s household income quartile decreased, stayed the same, or increased, respectively. We then examine whether income changes are associated with changes in fair-mindedness  $\hat{\alpha}_n$  and/or equality-efficiency orientation  $\hat{\rho}_n$ .

These results are presented in Table 5. In columns (1) and (5), we include the baseline 2013 estimates of  $\hat{\alpha}_n$  (column (1)) and  $\hat{\rho}_n$  (column (5)), the baseline income quartile measures, and the change in income quartile. In columns (2) and (6), we additionally include a rich set of covariates, including individual demographic controls. (To economize on space, we suppress the coefficients on control variables in Table 5, and provide the full regression output in Appendix Table A2.) In columns (1) and (2) we observe a positive and significant association between income gains and self-interest ( $p = 0.004$  in column (2)); we observe no significant association between income changes and efficiency-equality orientation in columns (5) and (6). These results are in line with the cross-sectional relationship between income and fair-mindedness – higher income is associated with greater self-interest – reported in earlier work (Fisman et al., 2017).

Finally, we look at whether political beliefs led to shifts in distributional preferences. We consider three possibilities: subjects that voted Democrat in the Presidential election in 2012 and in 2016; those that voted Republican in both elections; and those that switched party. To this end, for each

subject, we define an indicator variable that is  $-1$  for those that shift to voting Republican,  $+1$  for those that shift to voting Democrat, and  $0$  for those that do not change, or for whom there is missing data on voting. These results appear in columns (3) and (4) for fair-mindedness  $\hat{\alpha}_n$  and (7) and (8) for equality-efficiency orientation  $\hat{\rho}_n$ . Perhaps intriguingly, we find that shifting political allegiances is not associated with a shift in distributional preferences. Indeed, the only significant correlate of 2016 distributional preferences is that stable Democrats become more equality-oriented.<sup>12</sup> We again obtain similar results if we also categorize the equality-efficiency orientation ( $\hat{\rho}_n$ ) of self-interested subjects ( $\hat{\alpha}_n = 1$ ). These results are presented in Appendix Table A3.

*[Table 5 here]*

## 6 Conclusion

We measure the distributional preferences of a large and diverse sample of Americans during 2013-2016 – a period of significant social and economic upheaval – by embedding modified dictator games that vary the relative price of redistribution in the American Life Panel (ALP). Because choices are from standard budget sets, we can use classical revealed preference analysis to assess whether subject behavior is consistent with rationality, and classical demand analysis to recover the underlying preferences by estimating constant elasticity of substitution (CES) utility functions at the level of the individual subject.

To our knowledge, this is the first paper examining the intertemporal stability of preferences (distributional or otherwise) across several years in a large-scale experiment with such a subject pool, using an experimental design that is sufficiently effective to capture overlap in preferences across time. We find that, despite the social and economic turmoil during this period, our individual-level estimates of fair-mindedness and equality-efficiency orientation based on choices made in 2013 are highly predictive of those estimates based on choices three years later. Large changes that do occur are predicted, to some degree, by shifts in income – higher incomes are associated with greater self-interest.

---

<sup>12</sup>While its significance is marginal and thus we do not wish to ascribe too strong an interpretation to this finding, it fits with the increased focus on redistribution in much of the Democratic Party, as exemplified by support for, say, Elizabeth Warren’s proposed wealth tax and her relative success as a presidential candidate, as well as the enthusiasm for Bernie Sanders’ ultimately unsuccessful 2020 presidential bid.

Philosophical theories of fairness – utilitarianism, for example, or Rawlsianism – require that, unlike ordinary preferences for own-consumption, distributional preferences are not mere tastes but instead take a certain time-invariant form. Our results show that people’s actual distributional preferences are very heterogeneous yet show stability over time. This has important policy implications. Standard cost-benefit analysis of public policies measures costs and benefits by reference exclusively to the narrow self-interests of the persons whom the policies being assessed will affect. But insofar as policies have distributive implications, an accurate cost-benefit analysis would have to include costs and benefits associated with distributional preferences.

Our findings similarly have implications for canonical political economy models, as voters with distributional preferences naturally incorporate the interests of others into how they cast their ballots. Indeed, as the Downs (1957) paradox makes familiar, a rational and purely self-interested person would not vote at all. Nevertheless, political economy models have tended to assume voters are entirely self-interested. It is a fruitful direction for further research to similarly extend median voter-style models to allow for voters who care about others’ welfare. The fact that voters have heterogeneous distributional preferences may make such a modeling exercise quite complicated, but the fact that preferences are stable over time suggests that it is meaningful to think about how a given population of distributional preferences may aggregate up to desired political platforms and ultimately translate into redistributive policy.

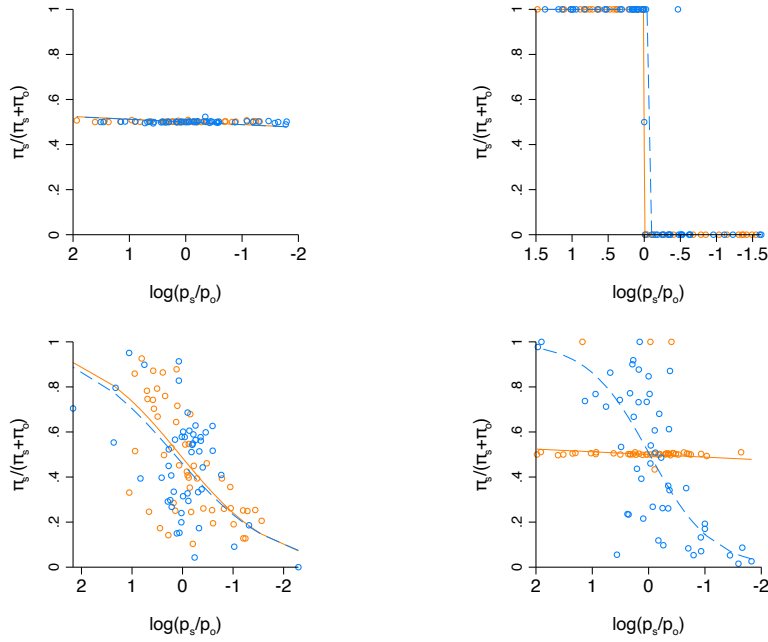
## References

- Afriat, Sydney N.**, “The Construction of Utility Functions from Expenditure Data,” *International Economic Review*, 1967, 8 (1), 67–77.
- , “Efficiency Estimates of Production Functions,” *International Economic Review*, 1972, 8, 568–598.
- Ahn, David, Syngjoo Choi, Douglas Gale, and Shachar Kariv**, “Estimating Ambiguity Aversion in a Portfolio Choice Experiment,” *Quantitative Economics*, 2014, 5 (2), 195–223.
- Alesina, Alberto, Stefanie Stantcheva, and Edoardo Teso**, “Intergenerational mobility and preferences for redistribution,” *American Economic Review*, 2018, 108 (2), 521–54.
- Almås, Ingvid, Alexander W Cappelen, and Bertil Tungodden**, “Cut-throat capitalism versus cuddly socialism: Are Americans more meritocratic and efficiency-seeking than Scandinavians?,” *Journal of Political Economy*, 2020, 128 (5), 1753–1788.
- Andreoni, James and John Miller**, “Giving according to GARP: An experimental test of the consistency of preferences for altruism,” *Econometrica*, 2002, 70 (2), 737–753.
- Becker, Gary S.**, “Irrational Behavior and Economic Theory,” *The Journal of Political Economy*, 1962, 70 (1), 1–13.
- Bellemare, Charles, Sabine Krger, and Arthur Van Soest**, “Measuring Inequity Aversion in a Heterogeneous Population Using Experimental Decisions and Subjective Probabilities,” *Econometrica*, 2008, 76 (4), 815–839.
- Bronars, Stephen G.**, “The Power of Nonparametric Tests of Preference Maximization,” *Econometrica*, 1987, 55 (3), 693–698.
- Bruhin, Adrian, Ernst Fehr, and Daniel Schunk**, “The many faces of human sociality: Uncovering the distribution and stability of social preferences,” *Journal of the European Economic Association*, 2019, 17 (4), 1025–1069.
- Camerer, Colin**, “Individual Decision Making,” in John Kagel and Alvin Roth, eds., *The Handbook of Experimental Economics*, Princeton: Princeton University Press, 1995, pp. 587–673.
- , *Behavioral Game Theory*, Princeton, NJ: Princeton University Press, 2003.
- Chambers, Christopher P. and Federico Echenique**, *Revealed Preference Theory* Econometric Society Monographs, Cambridge University Press, 2016.

- Charness, Gary and Matthew Rabin**, “Understanding social preferences with simple tests,” *The Quarterly Journal of Economics*, 2002, *117* (3), 817–869.
- Choi, Syngjoo, Raymond Fisman, Douglas Gale, and Shachar Kariv**, “Consistency and Heterogeneity of Individual Behavior under Uncertainty,” *American Economic Review*, 2007, *97* (5), 1921–1938.
- , **Shachar Kariv, Wieland Müller, and Dan Silverman**, “Who Is (More) Rational?,” *American Economic Review*, 2014, *104* (6), 1518–1550.
- Chuang, Yating and Laura Schechter**, “Stability of experimental and survey measures of risk, time, and social preferences: A review and some new results,” *Journal of Development Economics*, 2015, *117*, 151–170.
- Cooper, David J and John H Kagel**, “Other-regarding preferences,” *The handbook of experimental economics*, 2016, *2*, 217.
- Diewert, W. Erwin**, “Afriat’s Theorem and some Extensions to Choice under Uncertainty\*,” *The Economic Journal*, 2012, *122* (560), 305–331.
- Doherty, Carroll, Jocelyn Kiley, and Bridget Jameson**, “Partisanship and political animosity in 2016,” *Pew Research Center*, 2016, *75*.
- Downs, A.**, *An Economic Theory of Democracy*, New York, NY: Harper, 1957.
- Fehr, Ernst and Klaus M Schmidt**, “A theory of fairness, competition, and cooperation,” *The quarterly journal of economics*, 1999, *114* (3), 817–868.
- Fisman, Raymond, Pamela Jakiela, and Shachar Kariv**, “How Did Distributional Preferences Change During the Great Recession?,” *Journal of Public Economics*, 2015, *128*, 84–95.
- , —, and —, “Distributional preferences and political behavior,” *Journal of Public Economics*, 2017, *155*, 1–10.
- , —, —, and **Daniel Markovits**, “The distributional preferences of an elite,” *Science*, 2015, *349* (6254), aab0096.
- , **Shachar Kariv, and Daniel Markovits**, “Individual Preferences for Giving,” *American Economic Review*, 2007, *97* (5), 1858–1876.
- Forsythe, Robert, Joel Horowitz, N. S. Savin, and Martin Sefton**, “Fairness in Simple Bargaining Games,” *Games and Economic Behavior*, 1994, *6* (3), 347–369.
- Harsanyi, John C.**, “Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility,” *Journal of Political Economy*, 1955, *63*, 309–321.

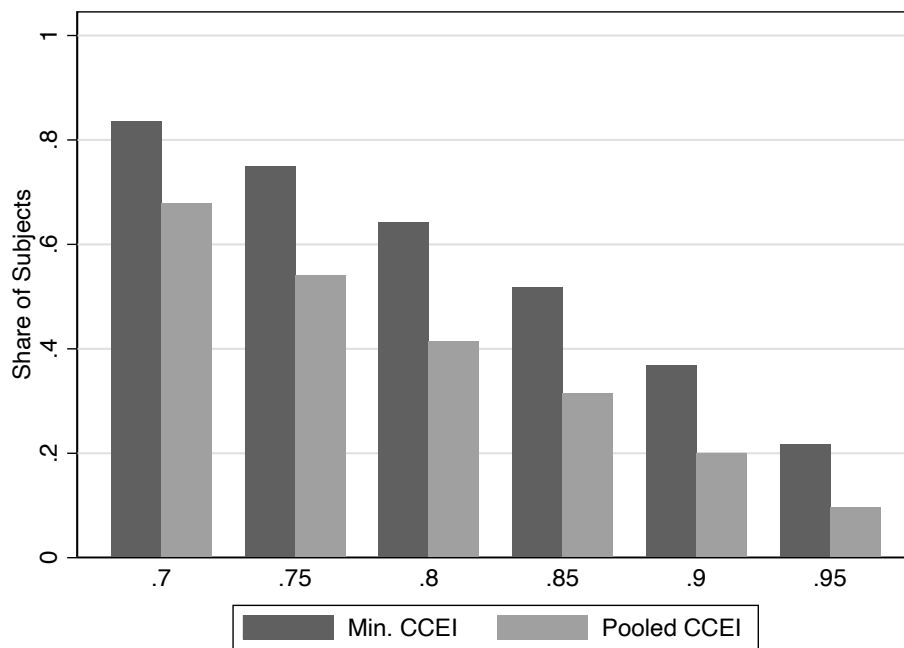
- , “Bayesian Decision Theory and Utilitarian Ethics,” *American Economic Review*, 1978, *68* (2), 223–228.
- Houtman, Martijn and J Maks**, “Determining all maximal data subsets consistent with revealed preference,” *Kwantitatieve methoden*, 1985, *19* (1), 89–104.
- Kuziemko, Ilyana, Michael I Norton, Emmanuel Saez, and Stefanie Stantcheva**, “How elastic are preferences for redistribution? Evidence from randomized survey experiments,” *American Economic Review*, 2015, *105* (4), 1478–1508.
- Levine, David K**, “Modeling altruism and spitefulness in experiments,” *Review of economic dynamics*, 1998, *1* (3), 593–622.
- Levitt, Steven D and John A List**, “What do laboratory experiments measuring social preferences reveal about the real world?,” *Journal of Economic perspectives*, 2007, *21* (2), 153–174.
- Li, Jing, William H Dow, and Shachar Kariv**, “Social preferences of future physicians,” *Proceedings of the National Academy of Sciences*, 2017, *114* (48), E10291–E10300.
- Rawls, John**, *A Theory of Justice*, Cambridge, MA: Harvard University Press, 1971.
- Schuerger, James M. and M. Tait E. Tavernelli**, “Temporal stability of personality by questionnaire,” *Journal of Personality and Social Psychology*, 1982, *43*, 176182.
- Varian, Hal R.**, “The Nonparametric Approach to Demand Analysis,” *Econometrica*, 1982, *50* (4), 945–972.
- , “Non-Parametric Tests of Consumer Behaviour,” *Review of Economic Studies*, 1983, *50* (1), 99–110.
- , “Goodness-of-Fit for Revealed Preference Tests,” working paper 1991.
- , “Revealed Preference and its Applications\*,” *The Economic Journal*, 2012, *122* (560), 332–338.
- Vermeulen, Frederic**, “FOUNDATIONS OF REVEALED PREFERENCE: INTRODUCTION,” *The Economic Journal*, 2012, *122* (560), 287–294.
- Zucman, Gabriel**, “Global wealth inequality,” *Annual Review of Economics*, 2019, *11*, 109–138.

Figure 1: The relationship between the log-price ratio ( $\log(p_s/p_o)$ ) and the token share ( $\pi_s/(\pi_s + \pi_o)$ ) for selected subjects.



*Note:* In each of the panels, orange circles indicate choices in 2013 and blue circles indicate choices in 2016. The solid orange line represents CES estimation in 2013 and a dotted blue line represents the CES estimation in 2016. These graphical comparisons provide some indication of goodness-of-fit for the four subjects we used for illustrative purposes. These figures are difficult to see in small black and white format. The scatterplots for the full set of subjects are available upon request.


Figure 2: The distributions of GARP violations



*Note:* The fraction of subjects whom the minimum CCEI for the separate datasets (black) and the CCEI for the combined dataset (gray) are above different critical values. The CCEI provides a summary statistic of the overall consistency of the data with GARP. The closer the CCEI 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.



Figure 3: The relationship between the distributional preferences in 2013 and 2016 at the individual-level



*Note:* Left panel: A scatterplots of the estimated fair-mindedness ( $\hat{\alpha}_n$ ) in 2013 and 2016. Middle panel: A scatterplots of the estimated equality-efficiency orientation ( $\hat{\rho}_n$ ) in 2013 and 2016. Right panel: the expenditure share of the tokens kept  $p_s \pi_s$  in 2013 and 2016. Blue triangles indicate statistically significant changes, while green circles indicate changes that are not statistically significant. The line in each panel is a polynomial regression plot with 95% confidence intervals. The histograms display the distribution of parameters in 2013 and 2016. These figures are difficult to see in small black and white format.

Table 1: Population-level summary of the individual-level CES estimates in 2013 and 2016

		2016				
		Selfish	Efficiency Focused	Intermediate	Equality Focused	Total
2013	Selfish	4.2	3.8	3.8	3.2	14.9
	Efficiency Focused	2.6	15.0	5.2	2.5	25.2
	Intermediate	4.6	9.8	15.9	7.4	37.6
	Equality Focused	2.6	3.3	6.2	10.2	22.3
Total		14.0	31.9	31.0	23.2	100

*Note:* We estimated individual-level CES parameters for a total of 687 subjects in both 2013 and 2016. We first classify subjects as purely self-interested ( $\hat{\alpha}_n = 1$ ) or fair-minded ( $\hat{\alpha}_n < 1$ ), and then classify the fair-minded subjects as equality-focused ( $\hat{\rho}_n < 0$ ), intermediate ( $\hat{\rho}_n = 0$ ) or efficiency-focused ( $\hat{\rho}_n > 0$ ). In our regression analysis below, we obtain near-identical results if we also categorize the equality-efficiency orientation ( $\hat{\rho}_n$ ) of self-interested subjects ( $\hat{\alpha}_n = 0$ ). These results are presented in Appendix Tables A1 and A2

Table 2: Comparing the ALP subjects with the U.S. population

	COMPLETED EXPERIMENT in 2013 (1)	COMPLETED EXPERIMENT in 2016 (2013) (2)	COMPLETED EXPERIMENT in 2016 (2016) (3)	ENTIRE ALP 2013 (4)	ENTIRE ALP 2016 (5)	US ADULTS 2012 (6)	US ADULTS 2016 (7)
Female	0.58	0.58	0.58	0.60	0.58	0.51	0.51
Age	49.37	49.55	52.36	49.05	52.96	46.68	47.29
18 to 44 years old	0.38	0.38	0.32	0.41	0.31	0.48	0.47
At least 65 years old	0.17	0.16	0.22	0.18	0.28	0.18	0.20
Non-Hispanic Caucasian	0.67	0.70	0.70	0.63	0.69	0.66	0.64
African American	0.11	0.10	0.10	0.12	0.11	0.12	0.12
Hispanic or Latino	0.18	0.16	0.16	0.21	0.16	0.15	0.16
High school diploma	0.91	0.92	0.92	0.93	0.96	0.88	0.89
College degree	0.31	0.34	0.34	0.36	0.44	0.27	0.29
HH Income (in tsd)	55.14	57.96	62.90	56.48	66.68	79.65	91.70
Currently employed	0.56	0.58	0.58	0.58	0.55	0.59	0.61
Currently unemployed	0.11	0.10	0.05	0.10	0.05	0.06	0.04
Out of labor force	0.34	0.32	0.37	0.32	0.39	0.35	0.36
Lives in northeast (census region I)	0.18	0.18	0.18	0.17	0.17	0.18	0.18
Lives in midwest (census region II)	0.20	0.21	0.21	0.19	0.19	0.21	0.21
Lives in south (census region III)	0.35	0.35	0.35	0.34	0.37	0.37	0.38
Lives in west (census region IV)	0.27	0.26	0.26	0.29	0.27	0.23	0.24

*Note:* Column (1) includes the 1,002 ALP respondents who completed the experiment in 2013. In columns (2) and (3), include the data for the subsample of 687 respondents who also completed the experiment in 2016 based on the 2013 and 2016 ALP questionnaire, respectively. Columns (4) and (5) compare our experimental subjects to the entire ALP sample in 2013 and 2016 and columns (6) and (7) compare them to the American Community Survey (ACS) conducted in 2012 and 2016. The ACS interviewed about 2.4 and 2.2 million respondents in 2012 and 2016, respectively. Averages are weighted to represent the adult population of the U.S. Place of residency is classified according to the Census Bureau regions: Northeast, Midwest, South and West.

Table 3: Comparing the household income of the ALP subjects and the U.S. Population

	COMPLETED EXPERIMENT in 2016 (2013) (1)	COMPLETED EXPERIMENT in 2016 (2016) (2)	US ADULTS 2012 (3)	US ADULTS 2016 (4)
Less than 10,000	8.76	8.47	5.20	4.44
10,000 to 19,999	13.0	11.2	8.36	6.82
20,000 to 29,999	11.1	11.4	9.28	7.91
30,000 to 39,999	9.93	11.4	9.40	8.40
40,000 to 49,999	11.4	10.1	8.93	8.14
50,000 to 59,999	10.1	6.42	8.21	7.77
60,000 to 74,999	10.9	11.5	10.9	10.6
75,000 to 99,999	9.20	8.91	13.6	14.0
100,000 to 124,999	5.99	9.05	9.22	10.1
125,000 to 199,999	7.30	8.03	11.3	13.8
200,000 and over	2.34	3.50	5.62	8.02

*Note:* Columns (1) and (2) report the household income of the subsample of 687 respondents who also completed the experiment in 2016 based on the 2013 and 2016 ALP questionnaire, respectively. Columns (3) and (4) report the household income in the American Community Survey (ACS) conducted in 2012 and 2016.

Table 4: The relationship between the distributional preferences in 2013 and 2016

	$\hat{\alpha}_n$ in 2016		$\hat{\rho}_n$ in 2016	
	(1)	(2)	(3)	(4)
$\hat{\alpha}_n$ in 2013	0.458*** (0.0356)	0.444*** (0.0377)		0.441 (1.742)
$\hat{\rho}_n$ in 2013		-0.00105 (0.00118)	0.444*** (0.0577)	0.434*** (0.0599)
CCEI in 2013		0.131** (0.0586)		-1.300 (1.942)
Observations	687	687	516	516
R-squared	0.203	0.213	0.203	0.204

*Note:* Robust standard errors in parentheses; \*, \*\*, and \*\*\* indicate 10, 5, and 1 percent significance levels, respectively. All specifications estimated via OLS. The dependent variable in columns (1)-(2) is fair-mindedness ( $\hat{\alpha}_n$ ) in 2016 and in columns (3)-(4) is equality-efficiency orientation ( $\hat{\rho}_n$ ) in 2016. The independent variables are the parameter estimates,  $\hat{\alpha}_n$  and  $\hat{\rho}_n$ , and the CCEI score in 2013. In columns (3) and (4), we report the results with  $\hat{\rho}_n$  in 2016 as the dependent variable for the 517 (75.3%) fair-minded subjects ( $\hat{\alpha}_n < 1$ ) in both 2013 and 2016. We obtain similar results if we also include the  $\hat{\rho}_n$  estimates of self-interested subjects ( $\hat{\alpha}_n = 0$ ). These results are presented in Appendix Table A1.

Table 5: The relationship between distributional preferences, economic circumstances and political preferences

	$\hat{\alpha}_n$ in 2016				$\hat{\rho}_n$ in 2016			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Changed Q	0.0296* (0.0161)	0.0353** (0.0140)			0.550 (0.596)	0.471 (0.574)		
Q2 in 2013	0.00232 (0.0204)	0.0211 (0.0177)			-1.577** (0.673)	-1.320** (0.631)		
Q3 in 2013	-0.00914 (0.0220)	0.0290 (0.0204)			-0.989 (0.980)	-1.153 (0.792)		
Q4 in 2013	0.0263 (0.0252)	0.0721*** (0.0254)			-0.924 (0.946)	-0.759 (0.996)		
Changed Party			0.0211 (0.0215)	-0.0112 (0.0187)			0.302 (0.861)	-0.224 (0.647)
Stayed Republican			-0.00747 (0.0197)	0.000765 (0.0174)			0.324 (0.830)	-0.430 (0.701)
Stayed Democrat			-0.0248 (0.0197)	-0.0297 (0.0181)			-1.239* (0.744)	-1.495** (0.686)
Other CES parameter in 2013	No	Yes	Yes	Yes	No	Yes	Yes	Yes
CCEI in 2013	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Individual demographics	No	Yes	No	Yes	No	Yes	No	Yes
Observations	687	684	687	684	516	513	516	513
R-squared	0.00633	0.287	0.00371	0.279	0.0113	0.307	0.00782	0.306

Robust standard errors in parentheses; \*, \*\*, and \*\*\* indicate 10, 5, and 1 percent significance levels, respectively. All specifications estimated via OLS. The dependent variable in columns (1)-(4) is fair-mindedness ( $\hat{\alpha}_n$ ) in 2016 and in columns (5)-(8) is equality-efficiency orientation ( $\hat{\rho}_n$ ) in 2016. The independent variable of interest in column (1)-(2) and (5)-(6) is an indicator variable which takes on values of  $-1, 0, 1$  based on whether the subject's household income quartile decreased, stayed the same, or increased, respectively. The independent variable of interest in column (3)-(4) and (7)-(8) is an indicator variable which takes on values of  $-1, 0, 1$  based on whether the subject shifted to voting Republican, did not change party (or has missing data on voting), or shifted to voting Democrat, respectively. Q2, Q3 and Q4 are indicator variables for the subject's household income quartile in 2013 (Q1 is the omitted category). In columns (5)-(8), we report the results with  $\hat{\rho}_n$  in 2016 as the dependent variable for the 517 (75.3%) fair-minded subjects ( $\hat{\alpha}_n < 1$ ) in both 2013 and 2016. We obtain similar results if we also include the  $\hat{\rho}_n$  estimates of self-interested subjects ( $\hat{\alpha}_n = 0$ ). These results are presented in Appendix Table A2. Individual demographic characteristics include gender, ethnicity, age and educational attainment. We provide the full regression output including individual demographic controls in Appendix Table A3.

## Appendix A

Table A1 : The relationship between the distributional preferences in 2013 and 2016, Full Sample

	$\hat{\alpha}_n$ in 2016		$\hat{\rho}_n$ in 2016	
	(1)	(2)	(3)	(4)
$\hat{\alpha}_n$ in 2013	0.458*** (0.0356)	0.444*** (0.0377)		1.176 (1.140)
$\hat{\rho}_n$ in 2013		-0.00105 (0.00118)	0.390*** (0.0515)	0.364*** (0.0516)
CCEI in 2013		0.131** (0.0586)		-3.628** (1.620)
Observations	687	687	687	687
R-squared	0.203	0.213	0.156	0.162

*Note:* The regression results reported in Table 4, also including the  $\hat{\rho}_n$  estimates of self-interested subjects ( $\hat{\alpha}_n = 0$ ) in columns (5)-(8). Robust standard errors in parentheses; \*, \*\*, and \*\*\* indicate 10, 5, and 1 percent significance levels, respectively. All specifications estimated via OLS. The dependent variable in columns (1)-(2) is fair-mindedness ( $\hat{\alpha}_n$ ) in 2016 and in columns (3)-(4) is equality-efficiency orientation ( $\hat{\rho}_n$ ) in 2016. The independent variables are the parameter estimates,  $\hat{\alpha}_n$  and  $\hat{\rho}_n$ , and the CCEI score in 2013.

Table A2 : The relationship between distributional preferences, economic circumstances and political preferences

	$\hat{\alpha}_n$ in 2016				$\hat{\rho}_n$ in 2016			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Changed Q	0.0296* (0.0161)	0.0353** (0.0140)			0.550 (0.596)	0.471 (0.574)		
Q2 in 2013	0.00232 (0.0204)	0.0211 (0.0177)			-1.577** (0.673)	-1.320** (0.631)		
Q3 in 2013	-0.00914 (0.0220)	0.0290 (0.0204)			-0.989 (0.980)	-1.153 (0.792)		
Q4 in 2013	0.0263 (0.0252)	0.0721*** (0.0254)			-0.924 (0.946)	-0.759 (0.996)		
$\hat{\alpha}_n$ in 2013		-0.565*** (0.0389)		-0.568*** (0.0386)		0.724 (1.803)		0.715 (1.796)
Female		0.00755 (0.0142)		0.000708 (0.0143)		-0.656 (0.547)		-0.647 (0.522)
Non-Hispanic Caucasian		-0.000573 (0.0160)		0.00669 (0.0164)		0.476 (0.598)		0.249 (0.581)
Q2 of Age		0.0188 (0.0191)		0.0224 (0.0195)		-1.749*** (0.643)		-1.598** (0.657)
Q3 of Age		0.0233 (0.0204)		0.0280 (0.0209)		-1.802*** (0.653)		-1.670** (0.671)
Q4 of Age		0.0320 (0.0201)		0.0287 (0.0206)		-1.562* (0.805)		-1.404* (0.785)
Some College		0.00640 (0.0177)		0.0169 (0.0176)		0.177 (0.621)		0.0874 (0.609)
College or more		-0.0347* (0.0181)		-0.000467 (0.0164)		-0.889 (0.696)		-0.960 (0.631)
$\hat{\rho}_n$ in 2013		-0.000645 (0.00128)		-0.000681 (0.00126)		-0.590*** (0.0613)		-0.597*** (0.0620)
CCEI in 2013		0.138** (0.0614)		0.136** (0.0605)		-1.669 (1.974)		-1.825 (2.005)
Changed Party			0.0211 (0.0215)	-0.0112 (0.0187)			0.302 (0.861)	-0.224 (0.647)
Stayed Republican			-0.00747 (0.0197)	0.000765 (0.0174)			0.324 (0.830)	-0.430 (0.701)
Stayed Democrat			-0.0248 (0.0197)	-0.0297 (0.0181)			-1.239* (0.744)	-1.495** (0.686)
Constant	-0.0248* (0.0146)	0.202*** (0.0540)	-0.0124 (0.0111)	0.223*** (0.0547)	1.072** (0.475)	1.795 (2.047)	0.504 (0.362)	1.771 (2.022)
Other CES parameter in 2013	No	Yes	Yes	Yes	No	Yes	Yes	Yes
CCEI in 2013	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Individual demographics	No	Yes	No	Yes	No	Yes	No	Yes
Observations	687	684	687	684	516	513	516	513
R-squared	0.00633	0.287	0.00371	0.279	0.0113	0.307	0.00782	0.306

*Note:* The regression results reported in Table 5, also including the  $\hat{\rho}_n$  estimates of self-interested subjects ( $\hat{\alpha}_n = 0$ ) in columns (5)-(8). Robust standard errors in parentheses; \*, \*\*, and \*\*\* indicate 10, 5, and 1 percent significance levels, respectively. All specifications estimated via OLS. The dependent variable in columns (1)-(4) is fair-mindedness ( $\hat{\alpha}_n$ ) in 2016 and in columns (5)-(8) is equality-efficiency orientation ( $\hat{\rho}_n$ ) in 2016. The independent variable of interest in column (1)-(2) and (5)-(6) is an indicator variable which takes on values of  $-1, 0, 1$  based on whether the subject's household income quartile decreased, stayed the same, or increased, respectively. The independent variable of interest in column (3)-(4) and (7)-(8) is an indicator variable which takes on values of  $-1, 0, 1$  based on whether the subject shifted to voting Republican, did not change party (or has missing data on voting), or shifted to voting Democrat, respectively. Q2, Q3 and Q4 are indicator variables for the subject's household income quartile in 2013 (Q1 is the omitted category).



Table A3 : The relationship between distributional preferences, economic circumstances and political preferences, Full Sample

	$\hat{\alpha}_n$ in 2016				$\hat{\rho}_n$ in 2016			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Changed Q	0.0296* (0.0161)	0.0353** (0.0140)			0.297 (0.525)	0.413 (0.485)		
Q2 in 2013	0.00232 (0.0204)	0.0211 (0.0177)			-1.259** (0.627)	-1.013* (0.563)		
Q3 in 2013	-0.00914 (0.0220)	0.0290 (0.0204)			-1.184 (0.820)	-1.091* (0.659)		
Q4 in 2013	0.0263 (0.0252)	0.0721*** (0.0254)			-0.691 (0.754)	-0.154 (0.785)		
Changed Party			0.0211 (0.0215)	-0.0112 (0.0187)			0.0450 (0.840)	-0.403 (0.618)
Stayed Republican			-0.00747 (0.0197)	0.000765 (0.0174)			0.127 (0.689)	-0.477 (0.574)
Stayed Democrat			-0.0248 (0.0197)	-0.0297 (0.0181)			-1.098* (0.656)	-1.314** (0.594)
Other CES parameter in 2013	No	Yes	Yes	Yes	No	Yes	Yes	Yes
CCEI in 2013	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Individual demographics	No	Yes	No	Yes	No	Yes	No	Yes
Observations	687	684	687	684	687	684	687	684
R-squared	0.00633	0.287	0.00371	0.279	0.00694	0.336	0.00477	0.336

A full regression output including individual demographic controls of the results presented in Table 5. Robust standard errors in parentheses; \*, \*\*, and \*\*\* indicate 10, 5, and 1 percent significance levels, respectively. All specifications estimated via OLS. The dependent variable in columns (1)-(4) is fair-mindedness ( $\hat{\alpha}_n$ ) in 2016 and in columns (5)-(8) is equality-efficiency orientation ( $\hat{\rho}_n$ ) in 2016. The independent variable of interest in column (1)-(2) and (5)-(6) is an indicator variable which takes on values of  $-1, 0, 1$  based on whether the subject's household income quartile decreased, stayed the same, or increased, respectively. The independent variable of interest in column (3)-(4) and (7)-(8) is an indicator variable which takes on values of  $-1, 0, 1$  based on whether the subject shifted to voting Republican, did not change party (or has missing data on voting), or shifted to voting Democrat, respectively. Q2, Q3 and Q4 are indicator variables for the subject's household income quartile in 2013 (Q1 is the omitted category). In columns (5)-(8), we report the results with  $\hat{\rho}_n$  in 2016 as the dependent variable for the 517 (75.3%) fair-minded subjects ( $\hat{\alpha}_n < 1$ ) in both 2013 and 2016. We obtain similar results if we also include the  $\hat{\rho}_n$  estimates of self-interested subjects ( $\hat{\alpha}_n = 0$ ).