

RESEARCH ARTICLE SUMMARY

ECONOMICS

The distributional preferences of an elite

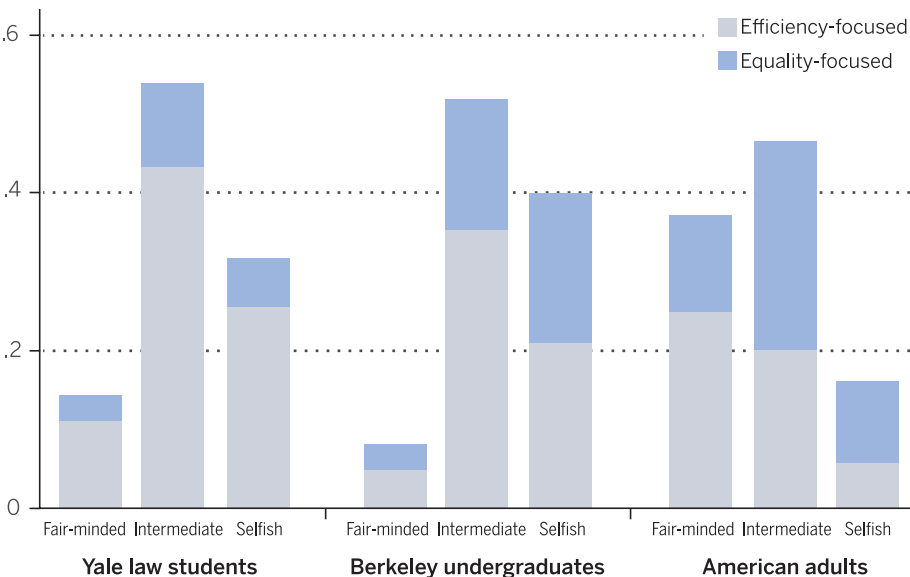
Raymond Fisman,* Pamela Jakiela, Shachar Kariv, Daniel Markovits

INTRODUCTION: Distributional preferences shape individual opinions and public policy concerning economic inequality and redistribution. We measured the distributional preferences of an elite cadre of Juris Doctor (J.D.) students at Yale Law School (YLS), a group that holds particular interest because they are likely to assume future positions of power and influence in American society. We compared the preferences of this highly elite group of students to those of a sample drawn from the American Life Panel (ALP), a broad cross-section of Americans, and to the preferences of an intermediate elite drawn from the student body at the University of California, Berkeley (UCB).

RATIONALE: We conducted modified dictator game experiments that varied the price of redistribution, i.e., the amount by which the “self’s” payoff must be decreased in order to increase the payoff of the “other” (an anonymous other subject) by one dollar. In contrast to

standard dictator games that do not vary the relative price of redistribution, our experimental design allows us to test whether our subjects’ preferences are formally rational and to decompose subjects’ preferences into two distinct tradeoffs: the tradeoff between self and other (fair-mindedness versus self-interest) and the tradeoff between equality and efficiency. For each subject, we estimated a constant elasticity of substitution (CES) utility function over payoffs to self and other; this functional form allows us to capture each tradeoff with a distinct parameter. A fair-minded subject places equal weight on the payoffs to self and other, whereas a selfish subject does not place any weight on the payoff to other; subjects’ preferences may also fall in between these two extremes. A subject with distributional preferences weighted toward equality (reducing differences in payoffs) increases the expenditure share spent on other as the price of redistribution increases, whereas a subject with distributional preferences weighted

Classifying subjects’ distributional preferences. We classify subjects as either fair-minded, intermediate, or selfish and as either equality-focused or efficiency-focused. The bars show the fraction of subjects in each category of self-interest in the elite YLS, UCB (the intermediate elite), and relatively less elite ALP samples. Each bar is then split into equality-focused and efficiency-focused subgroups, denoted by blue and gray, respectively.



toward efficiency (increasing total payoffs) decreases the expenditure share spent on other as the price of redistribution increases. An important strength of our measure of equality-efficiency tradeoffs between self and other is that it has been shown to predict such trade-

ON OUR WEB SITE

Read the full article at <http://dx.doi.org/10.1126/science.aab0096>

offs in distributional settings involving multiple others and to predict the likelihood of voting for political candidates perceived as favoring greater government redistribution.

This work therefore captures, in an experimental setting, a plausible measure of subjects’ attitudes toward actual redistributive policies.

RESULTS: YLS subjects were substantially more efficiency-focused than were the ALP subjects drawn from the general population. Overall, 79.8% of YLS subjects were efficiency-focused, versus only 49.8% of the ALP sample. The YLS subjects displayed this distinctive preference for efficiency over equality in spite of overwhelmingly (by more than 10 to 1) self-identifying as Democrats rather than Republicans. In addition, YLS subjects were less likely to be classified as fair-minded and more likely to be classified as selfish than were the ALP subjects. Subjects from the intermediate elite fell between the YLS and ALP subjects with respect to efficiency-mindedness but were less likely to be fair-minded and more likely to be selfish than were the YLS subjects. We also demonstrate the predictive validity of our experimental measure of equality-efficiency tradeoffs by showing that it predicts the subsequent career choices of YLS subjects: More efficiency-focused behavior in the laboratory was associated with a greater likelihood of choosing private sector employment after graduation, whereas more equality-focused behavior was associated with a greater likelihood of choosing nonprofit sector employment.

CONCLUSION: Our findings indicate sharp differences in distributional preferences between subjects of varying degrees of eliteness. These results provide a starting point for future research on the distinct preferences of the elite and differences in distributional preferences across groups more generally. From a policy perspective, our results suggest a new explanation for the modesty of the policy response to the rise in income inequality in the United States: Regardless of party, the policymaking elite is significantly more focused on efficiency vis-a-vis equality than is the U.S. public. ■

The list of author affiliations is available in the full article online.
*Corresponding author. E-mail: rfisman@bu.edu
Cite this article as R. Fisman et al., *Science* 349, aab0096 (2015). DOI: 10.1126/science.aab0096

RESEARCH ARTICLE

ECONOMICS

The distributional preferences of an elite

Raymond Fisman,^{1*} Pamela Jakiela,² Shachar Kariv,³ Daniel Markovits⁴

We studied the distributional preferences of an elite cadre of Yale Law School students, a group that will assume positions of power in U.S. society. Our experimental design allows us to test whether redistributive decisions are consistent with utility maximization and to decompose underlying preferences into two qualitatively different tradeoffs: fair-mindedness versus self-interest, and equality versus efficiency. Yale Law School subjects are more consistent than subjects drawn from the American Life Panel, a diverse sample of Americans. Relative to the American Life Panel, Yale Law School subjects are also less fair-minded and substantially more efficiency-focused. We further show that our measure of equality-efficiency tradeoffs predicts Yale Law School students' career choices: Equality-minded subjects are more likely to be employed at nonprofit organizations.

Growing economic inequality has intensified interest in the distinctive attitudes and behaviors of the American elite, whose sense of entitlement increasingly captures both general and scholarly attention (1). The interest in elites is not just voyeuristic, but practical; elites, and in particular graduates of elite universities and professional schools, exert considerable influence over public and private policy in the United States today. For example, over the past century more than half of the presidents, including the past four, attended Yale, Harvard, or Princeton. The preferences of a relatively small number of current students will therefore have a large and highly disproportionate impact on the future of the country as a whole.

We studied the distributional preferences of an important pool of future elite policy-makers and citizens: Juris Doctor (J.D.) students at the Yale Law School (YLS). As Alexis de Tocqueville observed in the 19th century, lawyers constitute an American aristocracy and wield an outsized influence over society in general and public policy in particular. Tocqueville's observations remain true today, particularly as they pertain to a very small number of top law schools, of which YLS is the most selective. Although the American Bar Association does not rank law schools, YLS has been ranked first in the country by U.S. News and World Report every year since 1987, when it began publishing the ranking (2). We compared the distributional preferences of this elite group of students to those of a sample drawn from the American Life Panel (ALP), an internet survey of a diverse population of U.S. adults.

Distributional preferences are important inputs into any measure of social welfare and enter every realm of policy-making. These preferences may naturally be divided into two qualitatively different components: the tradeoff between fair-mindedness and self-interest and the tradeoff between equality and efficiency. Although these two components of distributional preferences often operate together, they are conceptually distinct.

First, policy-makers must constantly decide whose interests matter and how much they matter. A baseline commitment to fair-mindedness—the ideal that all persons' interests matter equally—should inform all legitimate public policy. But in practice, fair-mindedness can be difficult to sustain against the many temptations to prefer one's own interests over the interests of others.

Second, policy-makers trade off equality and efficiency because reducing economic inequality almost inevitably has a cost; to use Okun's (3) famous metaphor, the transfer mechanisms that promote equality all involve leaky buckets. Policy-makers must thus decide, both in general and in any number of particular cases, by how much they are prepared to reduce aggregate income in order to secure a more equal income distribution. A comparison of the familiar philosophical theories of distributive justice—utilitarianism, for example, and Rawlsianism—further emphasizes the reasonable disagreements that fair-minded (impartial) policy-makers may have in trading off equality and efficiency.

In order to study the distributional preferences of an elite, we conducted laboratory experiments with YLS students using modified dictator games that vary the relative price of redistribution, building on the experiment first used by Andreoni and Miller (4). These decision problems are presented by using a graphical experimental interface that allows for the collection of rich individual-level data sets, as in (5). Specifically, we study a dictator game in which a

subject divides an endowment between “self” and an anonymous “other.” We denote persons self and other by s and o , respectively, and the associated monetary payoffs by π_s and π_o . In each decision problem, self allocates a unit endowment to π_s and π_o at fixed price levels p_s and p_o so that $p_s\pi_s + p_o\pi_o = 1$. This configuration creates budget sets over π_s and π_o in which $p = p_o/p_s$ is the relative price of redistribution.

The choice from a budget set indicates a subject's preferred allocation relative to a broad range of possible alternatives; it therefore provides more information about preferences than a choice from a discrete set of options would reveal. Furthermore, because of the user-friendly experimental interface, it is possible to present each subject with many choices in the course of a single experimental session, yielding an extremely rich data set. These data allow us to apply powerful techniques from demand analysis to determine whether each subject's behavior is consistent with utility maximization and to identify the structure of the utility function that rationalizes each subject's choices.

Our analysis examines the differences between the distributional preferences of elite YLS subjects—particularly, their willingness to sacrifice efficiency to reduce inequality—and the distributional preferences of the diverse sample of (relatively less elite) Americans in the ALP subject pool. In contrast to standard dictator games that do not vary the relative price of redistribution, our design allows us to separate fair-mindedness from equality-efficiency tradeoffs by examining subjects' responses to changes in the relative price of redistribution. A subject who decreases the expenditure share spent on other, $p_o\pi_o$, when the relative price of redistribution p increases has distributional preferences weighted toward efficiency (increasing total payoffs), whereas a subject who increases $p_o\pi_o$ when p increases has distributional preferences weighted toward equality (reducing differences in payoffs). For each subject, we constructed a measure of equality-efficiency tradeoffs by estimating a constant elasticity of substitution (CES) utility function over payoffs to self and other. A strength of this measure is that it has been shown to predict the equality-efficiency tradeoffs in distributional settings involving multiple others (5) and to predict the likelihood of voting for political candidates perceived as favoring greater government redistribution (6). We further validated the external validity of our measure in the present study by showing that it predicts YLS subjects' subsequent career choices. Taken together, this suggests that our measure of equality-efficiency tradeoffs meaningfully captures individual distributional preferences that govern subjects' real-world decisions.

Subject pools YLS subjects

We conducted experimental sessions at YLS during the spring semesters of 2007, 2010, and 2013. The 3-year lag between experiments means that each set of sessions draws from an entirely

¹Department of Economics, Boston University, Boston, MA, USA.

²Department of Agricultural and Resource Economics, University of Maryland, College Park, MD, USA. ³Department of Economics, University of California, Berkeley, Berkeley, CA, USA. ⁴Yale Law School, Yale University, New Haven, CT, USA.

*Corresponding author. E-mail: rfisman@bu.edu

new YLS student body. Of the 208 subjects in the YLS sample, 199 reported their year of study; 91 subjects were 1st-year students, 61 were 2nd-year students, and the remainder were 3rd-year students. Summary statistics on the basic sociodemographic characteristics of our two main pools of subjects, the YLS and ALP samples, are reported in Table 1 (7).

YLS enrolls about 200 students per year, making it among the smallest and most selective graduate law schools in the United States. In the most recent year for which data are available, YLS accepted only 11.3% of its (already elite) college-educated applicants. YLS students tend to come from educated, relatively well-off households. In our experiments, 95 YLS subjects reported that both parents hold graduate degrees, and 113 grew up in U.S. ZIP codes where the average household income was above \$70,000 in 2014 inflation-adjusted dollars (the mean household income in the U.S. was \$72,641 in 2014). YLS students also have extremely high expected future incomes; although YLS does not disclose the starting salaries of its graduates, the median starting salary for graduates at top law schools such as Yale, Harvard, and Columbia is \$160,000 per year (often augmented by signing bonuses). Overall, the YLS subjects are one of the most academically elite groups in the United States and can, in expectation, expect to join the ranks of the economic and political elite as well.

ALP subjects

For comparative purposes, we present our YLS data alongside a subset of the data of (6), collected in 2013 by using near-identical experiments with the ALP, an internet survey of more than 5000 adult Americans. The overall sample of ALP respondents is broadly comparable with the U.S. population in terms of demographic and socioeconomic characteristics; it includes an enormous amount of demographic, socioeconomic, and geographic diversity. Fisman *et al.* provide a detailed comparison of ALP subjects with respondents from the American Community Survey (ACS) conducted by the U.S. Census and representative of the U.S. population (6). The subsample of 1002 ALP respondents in the subject pool described in (6) is remarkably consistent with the entire ALP sample.

To focus on ALP subjects comparable in age with YLS subjects, we restricted attention to the

309 subjects in the original sample who were aged 40 and under. Summary statistics on the basic sociodemographics of the 309 subjects included in our analysis are reported in Table 1. As Table 1 indicates, the overwhelming majority of the ALP subjects are less educated than the YLS subjects.

Intermediate elites

We probed the generalizability of our results with the YLS and ALP samples by examining the behavior of two intermediate elites. This can help to rule out, in particular, the possibility that we are simply picking up a law school effect. First, we compared the most highly educated, wealthy ALP subjects to a nonelite comparison group of ALP subjects with less education and income. Second, we compared subjects drawn from the large and diverse student body of University of California, Berkeley (UCB) undergraduates to the full sample of ALP subjects (aged 40 and under). By examining these two intermediate elites, including one drawn from a broad cross-section of the general (primarily nonstudent) population, our aim was to assess the extent to which our conclusions are likely to reflect the distinctive distributional preferences of the U.S. elite, and not just those of YLS students or, more broadly, those in the legal profession.

ALP elite

We classify an ALP respondent as elite if she or he is employed, reported an annual household income over \$100,000, and holds a graduate degree. In the experiments of Fisman *et al.*, only 9 of the 309 subjects aged 40 and under met these criteria (6). To obtain a larger sample of elite ALP subjects, we conducted an additional round of experiments in 2014, inviting all ALP respondents who met our criteria for elitiness to participate, along with a comparison group of nonelite ALP respondents (who were also employed and aged 40 and under but reported household incomes below \$100,000 and did not hold graduate degrees). Combining the data on elite and nonelite ALP subjects from our two waves of experiments, we defined two additional samples: ALP elite (54 subjects) and ALP nonelite (206 subjects) (8).

UCB student elite

We also examined a second intermediate elite: undergraduate students at UCB, which is ranked

as the world's top public university and among the most prestigious universities—public or private—globally. As a large public university, UCB draws its students from a diverse range of socioeconomic and cultural backgrounds. It is therefore a useful additional comparator for assessing both whether the patterns we attributed to the YLS subjects' elitiness hold more broadly, and also for emphasizing the extreme preferences we observed in the highly elite YLS sample. To this end, we used data collected in 2004 and 2011 in identical experiments at the UCB Experimental Social Science Laboratory (Xlab). The Xlab draws its subjects from all students and administrative staff, but most subjects in its experiments are undergraduate students. Fisman *et al.* describe the make-up of UCB student population during 2004–2011 and the composition of the UCB subjects in these experiments (9).

The experiment

In our experiments, we presented subjects with a sequence of 50 decision problems in which each choice has consequences for self (the subject) and for an anonymous other. Each decision problem is presented as a choice from a two-dimensional budget set. A choice of the allocation (x, y) from the budget set represents an allocation between accounts: Self received the tokens allocated to the y account, and other received the tokens allocated to the x account. More precisely, each decision involved choosing a point on a budget set over possible token allocations to self and other so that $p_x \pi_s + p_y \pi_o = 1$, where π_s and π_o correspond to the payoffs to self and other, respectively, and $p = p_o/p_s$ is the relative price of redistribution. In each decision problem, the computer selected a budget set at random from the set of budget sets that intersected at least one of the axes at 50 or more experimental tokens, but with neither intercept exceeding 100 tokens. These decision problems were presented by use of a graphical interface, and choices were made by using the mouse to move the pointer on the screen to the desired point (10). At the end of the experiment, one of each subject's choices was randomly selected to determine final payouts.

Framework for analysis Nonparametric analysis

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 provides a direct test (11–13): Choices from a finite collection of budget sets are consistent with maximizing a piecewise linear, continuous, increasing, and concave utility function if and only if they satisfy the Generalized Axiom of Revealed Preference (GARP) (14). Hence, to assess whether an individual subject's choice data are consistent with utility-maximizing behavior, we needed to check whether the data satisfy GARP. Because our subjects make choices over a wide range of intersecting budget sets, our data provide a stringent test of utility maximization (15).

Table 1. Summary statistics on subjects in YLS and ALP samples.

	Subject pool	
	YLS subjects	ALP subjects
Age	25.40	31.23
Female	0.466	0.653
Born in the United States	0.782	0.906
Non-Hispanic white	0.632	0.545
Completed college	1	0.327
Observations	208	309




Fig. 1. CCEIs. (A and B) Histograms of the CCEI in (A) the YLS and ALP samples and (B) the ALP elite versus nonelite samples. CCEIs closer to 1 mean the data are closer to perfect consistency with GARP and hence to perfect consistency with utility maximization.

GARP provides a discrete test of utility maximization—either the data satisfy GARP or they do not—but individual choices frequently involve errors; subjects may compute incorrectly, execute intended choices incorrectly, or err in other less obvious ways. To account for the possibility of errors, we assessed how nearly individual choice behavior complies with GARP by using Afriat’s (16) 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 between 0 and 1: indices closer to 1 mean that the data are closer to perfect consistency with GARP and hence to consistency with utility maximization.

Parametric analysis

In the case of two goods, consistency with GARP and budget balancedness together imply that the demand function is homogeneous of degree zero. If we also assume separability and homotheticity, then the underlying utility function must be a member of the constant elasticity of substitution (CES) family (17). The CES utility function is given by

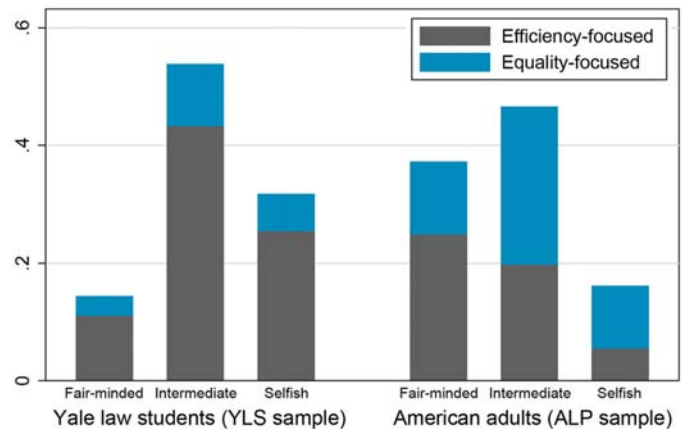
$$u_s(\pi_s, \pi_o) = [\alpha(\pi_s)^\rho + (1-\alpha)(\pi_o)^\rho]^{1/\rho}$$

where α represents the relative weight on the payoff for self vis-a-vis other (fair-mindedness), and ρ represents the curvature of the indifference curves (equality-efficiency tradeoffs).

Those with $\alpha = 1/2$ are fair-minded in the sense that they place equal weight on the payoffs to self and other; those with $\alpha = 1$ are perfectly selfish and do not put any weight on the payoff to other. Those with $1/2 < \alpha < 1$ exhibit some (intermediate) degree of fair-mindedness. For any $\rho > 0$, an increase in the relative price of redistribution raises—and for any $\rho < 0$, and increase in the relative price of redistribution lowers—the expenditure share on tokens allocated to self, $p_s \pi_s$. Thus, those with $\rho > 0$ have distributional preferences weighted toward maximizing efficiency (increasing total payoffs),

Fig. 2. Classifying subjects’ distributional preferences.

We classify subjects as either fair-minded ($\hat{\alpha}_n < 0.95$), or selfish ($\hat{\alpha}_n > 0.95$) and as either equality-focused ($\hat{\rho}_n < 0$) or efficiency-focused ($\hat{\rho}_n > 0$). The bars show the fraction of subjects in each category of self-interest in the YLS and ALP samples. Each bar is then split into equality-focused and efficiency-focused subgroups, denoted by blue and gray, respectively.



whereas those with $\rho < 0$ have distributional preferences weighted toward minimizing inequality (reducing differences in payoffs).

Our estimation was done for each subject n separately, generating individual-level estimates of the CES parameters. Specifically, we normalized prices at each observation and estimated demand in terms of budget shares, which are bounded between 0 and 1, using nonlinear Tobit maximum likelihood.

Experimental results

In this section, we provide results from the YLS and ALP samples. We first examine whether the data observed in our experiment could have been generated by a subject maximizing a well-behaved utility function. We then proceed to our econometric analysis by imposing further structure on the data in order to recover the underlying distributional preferences.

Rationality

The mean CCEI in the YLS sample is 0.95, and the median is 0.99, indicating that the overwhelming majority of the YLS subjects make choices that are perfectly or almost perfectly

consistent with utility maximization. For comparison, the mean CCEI in the ALP sample is 0.86, and the median is 0.89. Thus, the choices of the ALP subjects are generally consistent with utility maximization (18). Nonetheless, the CCEIs of the YLS subjects are substantially higher than those of subjects in the ALP sample. A Wilcoxon rank-sum test rejects the hypothesis that the distributions of CCEIs are equal ($P < 0.001$). Histograms of CCEIs of both the YLS and ALP subjects are presented in Fig. 1A. Relative to the CCEIs of the general population in the ALP sample, the CCEIs of YLS subjects are skewed to the right (19).

Preferences

Our subjects’ CCEIs are sufficiently near 1 to justify treating the data as utility-generated. If we also assume separability and homotheticity, then the underlying utility function $u_s(\pi_s, \pi_o)$ that rationalizes the data must be a member of the CES family. We now turn to the analysis of our estimates of the individual CES utility parameters, $\hat{\alpha}_n$ and $\hat{\rho}_n$. The distributions of $\hat{\alpha}_n$ and $\hat{\rho}_n$ in the YLS and the ALP samples are summarized in Fig. 2. Across all categories of self-interest (fair-minded, intermediate, and selfish), the YLS

subjects are substantially more efficiency-focused than are the ALP subjects drawn from the general population. Overall, 79.8% of YLS subjects are efficiency-focused ($\hat{\rho}_n > 0$) versus only 49.8% of the ALP sample. In addition, the YLS subjects are less likely to be classified as fair-minded and more likely to be classified as selfish than are the ALP subjects: 14.4% of YLS subjects are classified as fair-minded, as compared with 37.2% of ALP subjects; conversely, 31.7 and 16.2% of YLS and ALP subjects, respectively, are classified as selfish (20).

The distributions of the parameter estimates $\hat{\alpha}_n$ and $\hat{\rho}_n$ for the both YLS and ALP samples are presented in Figs. 3 and 4, respectively. The distribution of the estimated $\hat{\alpha}_n$ parameters of the YLS sample is skewed sharply to the right, relative to the distribution of the ALP sample (Fig. 3A); the YLS subjects are substantially less likely to have estimated $\hat{\alpha}_n$ parameters below 0.6 (22.1 versus 52.4%), are somewhat more likely to have estimated $\hat{\alpha}_n$ parameters between 0.6 and 0.9 (42.3 versus 27.5 %), and substantially more likely to have estimated $\hat{\alpha}_n$ parameters above 0.9 (35.6 versus 20.1%). As shown in Fig. 4A, the distribution of the estimated $\hat{\rho}_n$ parameters of the YLS sample lies clearly to the

right of the ALP sample's distribution of $\hat{\rho}_n$ values, indicating a much higher degree of efficiency orientation in our elite sample. A substantial majority of YLS subjects have estimated $\hat{\rho}_n$ parameters above 0.5 (60.1% of subjects), indicating a very high degree of efficiency focus; for comparison, only 17.8% of ALP subjects have estimated $\hat{\rho}_n$ parameters that high.

We next turned to regression analyses that more systematically examine the differences in fair-mindedness ($\hat{\alpha}_n$) and equality-efficiency tradeoffs ($\hat{\rho}_n$) between the YLS and ALP samples. We defined an indicator variable to denote the YLS sample and present the results of individual-level regressions with this as the primary explanatory variable in Table 2, which includes the results with no individual-level controls and when we control for gender, age, and education level (having a college degree).

In the first column of Table 2, we present a specification with the fair-mindedness parameter as the dependent variable, using a Tobit model that allows for the censoring of $\hat{\alpha}_n$ at 1. The parameters are, on average, 0.12 higher in the YLS sample than in the ALP sample ($P < 0.001$), indicating that the YLS subjects are substantially and significantly less fair-minded than the ALP

subjects. After controlling for age, gender, and education level, we still find that YLS subjects are significantly less fair-minded than are ALP subjects ($P = 0.002$), but the point estimate drops to 0.084.

We next present a specification with the equality-efficiency tradeoff parameter as the dependent variable. Because the distribution of $\hat{\rho}_n$ is highly skewed ($\hat{\rho}_n$ ranges from $-\infty$ to 1 and a number of subjects have very low values of $\hat{\rho}_n$), we estimate quantile regressions that are less sensitive to extreme values. We report results for the 25th, 50th, and 75th percentiles in second through fourth columns of Table 2, respectively. Across all three quantiles, the values of $\hat{\rho}_n$ are significantly higher in the YLS sample than in the ALP sample ($P < 0.001$), indicating that the YLS subjects are significantly more focused on efficiency vis-a-vis equality than are the ALP subjects. After adding controls, the point estimates on YLS are reduced by about one half, but for both the 50th and 75th percentiles, the coefficient remains significant ($P = 0.003$ and 0.001, respectively).

As an alternative approach to dealing with the skewed distribution of $\hat{\rho}_n$, in the fifth column of Table 2 we present a probit specification using an

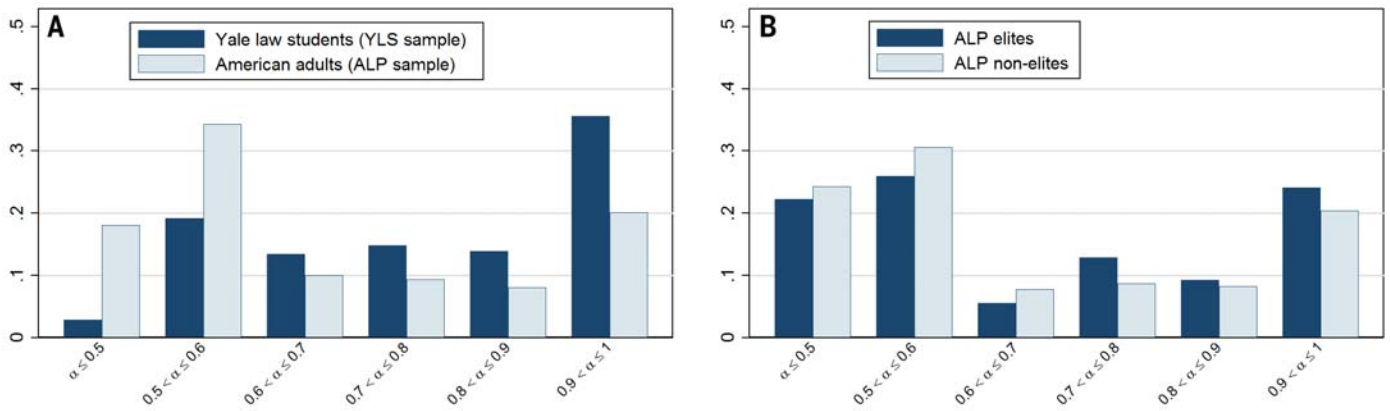


Fig. 3. Estimated $\hat{\alpha}_n$ parameters. (A and B) Histograms of the $\hat{\alpha}_n$ estimates in (A) the YLS and ALP samples and (B) the ALP elite versus nonelite samples. $\hat{\alpha}_n$ indexes fair-mindedness: the relative utility weight placed on one's own payoff vis-à-vis the payoff to other.

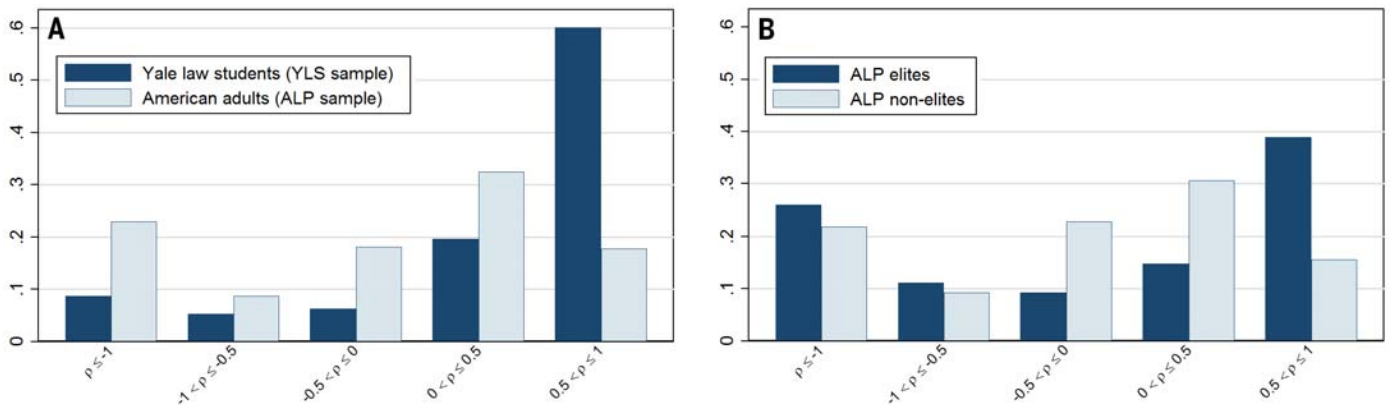


Fig. 4. Estimated $\hat{\rho}_n$ parameters. (A and B) Histograms of the $\hat{\rho}_n$ estimates in (A) the YLS and ALP samples and (B) the ALP elite versus nonelite samples. $\hat{\rho}_n$ indexes equality-efficiency tradeoffs; $\hat{\rho}_n$ values closer to 1 indicate greater efficiency focus.

Table 2. Regressions of estimated CES parameters, by subject pool. Standard errors are in parentheses. ***, significance at the 99% level; **, significance at the 95% level; *, significance at the 90% level.

Specification	Quantile regressions				
	Tobit	25th percentile	50th percentile	75th percentile	Probit
Dependent variable	$\hat{\alpha}_n$	$\hat{\rho}_n$	$\hat{\rho}_n$	$\hat{\rho}_n$	$I(\hat{\rho}_n > 0)$
	<i>Without controls</i>				
YLS student	0.120*** (0.017)	1.075*** (0.217)	0.693*** (0.080)	0.405*** (0.049)	0.831*** (0.122)
Constant	0.668*** (0.011)	-0.867*** (0.137)	0.005 (0.051)	0.420*** (0.031)	0.004 (0.071)
Observations	517	517	517	517	517
	<i>Including controls for age, gender, and education level</i>				
YLS student	0.084*** (0.027)	0.552 (0.476)	0.357*** (0.120)	0.260*** (0.077)	0.440** (0.183)
Constant	0.660*** (0.029)	-0.051 (0.480)	0.260** (0.121)	0.604*** (0.078)	0.497*** (0.183)
Observations	514	514	514	514	514

indicator for efficiency-oriented subjects ($\hat{\rho}_n > 0$) as the dependent variable. We found that the YLS subjects are 29.2 percentage points more likely to be efficiency-focused than are the ALP subjects ($P < 0.001$). After controlling for demographics, the YLS subjects are still 14.1 percentage points more likely to be efficiency-focused than are the ALP subjects ($P = 0.016$).

Our results are thus robust to the inclusion of controls for age, gender, and education. Education is a defining feature of the elite, and as such, whether it should be accounted independently for its role is unclear. Still, we argue in the spirit of Altonji *et al.* (21) that if unobserved attributes—which we would have expected a priori to be correlated with education—were a dominant source of the observed correlation between elite status and distributional preferences, then adding controls should have had a substantial effect on the estimated associations.

External validity—YLS subjects' career choices

Our results above show that subjects drawn from the student population at YLS—the future U.S. elite—are much more rational (in the sense of implementing a consistent, complete, and transitive preference ordering) and are far more inclined to favor efficiency over equality relative to subjects drawn from the ALP, a diverse cross-section of Americans. Yet, this analysis rests on the assumption of external validity; we assume that our individual-level laboratory experimental measure of equality-efficiency tradeoffs predicts the willingness to trade off equality and efficiency outside the laboratory. As discussed above, our experimental design was selected in part because it has been shown to predict equality-efficiency tradeoffs in a range of experimental settings (5) and to predict voting behavior (6). To further assess the external validity of our experimental measure of equality-efficiency tradeoffs, we tested whether YLS subjects' dis-

tributional preferences, as captured in our experiment, are reflected in behavior in a natural decision environments by looking at subjects' (early) career choices.

In late 2014, we obtained approval to access the names of subjects in the first two waves of the YLS experiment fielded in 2007 and 2010 (subjects who participated in 2013 are still students at YLS or at extremely early stages of their careers). We were able to track down, via Web searches, the career choices of 137 out of the 139 subjects (22). Of these, 119 subjects (86.9%) could be cleanly classified based on employer type: nonprofit (33 subjects), academia (13 subjects), government (18 subjects), and corporate (66 subjects). Of the remaining 17 subjects, 14 subjects extended their training as clerks, a position that can serve as preparation for a range of legal careers, and three continued their schooling.

YLS graduates who chose nonprofit careers tended to pursue the equality-related rights and interests of the disenfranchised. In contrast, YLS graduates who work in the corporate sector overwhelmingly serve as managers or deal-makers whose basic purpose is to extract efficiencies on behalf of their employers or clients. Moreover, although this observation is perhaps more anecdotal, the corporate workplace itself is more single-mindedly structured around efficiency than are workplaces in the nonprofit sector. We grouped government and academia as an intermediate case and examined whether the nonprofit and corporate subsamples have substantially and significantly lower and higher, respectively, efficiency orientations relative to other YLS subjects.

We assert that the existence of a relationship between our experimental measure of the equality-efficiency tradeoffs of YLS subjects and their real-world career choices would confer substantial external validity on the conclusions drawn from our laboratory experiments (23). The median $\hat{\rho}_n$ parameter value among YLS subjects employed in the nonprofit, academia/government, and cor-

porate sectors are 0.439, 0.648 and 0.745, respectively. A Wilcoxon rank-sum test rejects the hypothesis that the nonprofit and corporate subsamples have equal $\hat{\rho}_n$ distributions ($P = 0.057$). A rank-sum tests rejects equality of the $\hat{\rho}_n$ distributions in the corporate and academia/government subsamples ($P = 0.035$) but (unsurprisingly given the small sample sizes) does not reject the hypothesis that the $\hat{\rho}_n$ distributions in the nonprofit and academia/government subsamples are equal ($P = 0.637$).

Last, we further investigated the relationship between the equality-efficiency tradeoffs of YLS subjects and their career choices using an ordered logit regression, ranking the (ascending) efficiency orientation of employment types as nonprofit, academia/government, or corporate. We report the results in Table 3. Given the skewed distribution of the estimated $\hat{\rho}_n$ parameters, we provide two alternative specifications: in Table 3, column 1, the independent variable is an indicator for having an above median (within the YLS sample) $\hat{\rho}_n$ value, whereas in column 2 it is the decile of the estimated $\hat{\rho}_n$ distribution. Table 3 also shows the results both with no individual-level controls and when we control for gender, age at the time the subject participated in the experiment, and the year of participation (either 2007 or 2010). The estimation results confirm our findings above in a regression setting: Our experimental measure of equality-efficiency tradeoffs predicts YLS subjects subsequent career choices.

Intermediate elites

We last examined whether our findings on the distinct distributional preferences of YLS subjects can plausibly be applied to elites more broadly. This analysis helps to ensure that we are not simply picking up an effect peculiar to the YLS population. To do so, we examined the behaviors of two intermediate elites who participated in identical experiments. First, we compared the most highly educated, wealthy ALP

Table 3. Ordered logit estimation of YLS subjects' career choices. Standard errors in parentheses. ***: significance at the 99% level; **: significance at the 95% level; *: significance at the 90% level. Dependent variable is equal to 1 for subjects who work in the nonprofit sector, equal to 2 for subjects who work in academia or government, and equal to 3 for subjects who work in the corporate sector. Controls are for age, gender, and year of experimental session.

Dependent variable: post-YLS career category

	Without controls	
Above median $\hat{\rho}_n$	1.043***	
	(0.364)	
Decile of estimated $\hat{\rho}_n$		0.157**
		(0.068)
Observations	120	120
	With controls	
Above median $\hat{\rho}_n$	1.035***	
	(0.374)	
Decile of estimated $\hat{\rho}_n$		0.164**
		(0.076)
Observations	118	118

respondents with a comparison group of non-elite ALP respondents with less education and income. Second, we compared UCB undergraduate students with the ALP respondents drawn from the general population. By examining two different intermediate elites, including one drawn from a broad cross-section of the general (primarily nonstudent) population, our aim was to assess the extent to which our conclusions are likely to reflect the distinctive distributional preferences of the U.S. elite, and not just those of YLS students.

Elite ALP subjects

We classify an ALP subject as elite if she or he is employed, has an annual household income of over \$100,000, and holds a graduate degree (although this definition does not approach the eliteness of the YLS subjects). We compared the elite ALP subjects with nonelite ALP subjects who are employed but with incomes below \$100,000 and no graduate degree. The ALP subsample we used for this elite versus nonelite analysis comprises data collected across two waves of experiments and includes 54 ALP elite and 206 ALP nonelite subjects. Paralleling our main analysis, histograms comparing CCEI scores, fair-mindedness ($\hat{\alpha}_n$), and equality-efficiency tradeoffs ($\hat{\rho}_n$) between the ALP elite and ALP nonelite subsamples are presented in Figs. 1B, 3B, and 4B, respectively.

As shown in Fig. 1B, the distribution of CCEI scores for the ALP elite subjects is skewed to the right relative to that of the ALP nonelites. The mean and median CCEI scores of ALP elites are 0.882 and 0.945 versus 0.857 and 0.891 for nonelites, and a Wilcoxon rank-sum test rejects the equality of the distributions ($P = 0.014$). A much more modest association between ALP elite status and fair-mindedness ($\hat{\alpha}_n$) is shown in Fig. 3B—the mean and median values are 0.687 and 0.612 versus 0.659 and 0.576 for elites and nonelites, respectively. This indicates, as with our main analysis, a lower level of fair-mindedness among the

elite, although here the difference in distributions is modest and statistically insignificant (Wilcoxon rank-sum test $P = 0.429$).

Last, there are sharp differences between the ALP elite and nonelite subjects in their equality-efficiency tradeoffs ($\hat{\rho}_n$) (Fig. 4B). Although we cannot reject the equality of the distributions (Wilcoxon rank-sum test $P = 0.209$), a much larger fraction of ALP elite subjects have high values of $\hat{\rho}_n$, indicating greater efficiency focus. Specifically, 38.9% of the ALP elite subjects have estimated $\hat{\rho}_n$ parameters of at least 0.5, as compared with only 15.5% for the ALP nonelites. In contrast, only 24.1% of the ALP elite subjects have intermediate $\hat{\rho}_n$ parameters between -0.5 and 0.5 as compared with 53.4% of ALP nonelite subjects. Thus, as with the YLS elite, we observed a stronger efficiency orientation among elites than among nonelites within the ALP subject pool (24).

Elite UCB subjects

We next turned to a second intermediate elite: undergraduates at UCB. Again paralleling our main analysis, histograms comparing the CCEIs, fair-mindedness ($\hat{\alpha}_n$), and equality-efficiency tradeoffs ($\hat{\rho}_n$) between the UCB subjects and the ALP subjects drawn from the general population are shown in Figs. S1, S2, and S3. Overall, the differences between the ALP subjects and the UCB subjects are very similar to the differences between the ALP and YLS subjects.

A sharp difference in CCEIs between the UCB subjects and ALP subjects is shown in Fig. S1, although the difference is slightly smaller than that observed in comparing the CCEIs of the YLS subjects with those of the ALP subjects. A Wilcoxon rank-sum test rejects the equality of the distributions ($P < 0.001$). Turning to the distributions of fair-mindedness ($\hat{\alpha}_n$), there is an even larger gap between the UCB and ALP subjects than between the YLS and the ALP subjects (reported in our main analysis); the median values for the UCB and ALP subject pools are 0.888

and 0.591, respectively (Wilcoxon rank-sum test $P < 0.001$) (fig. S2).

Last, there is a sharp difference between the UCB and ALP subjects in their equality-efficiency tradeoffs ($\hat{\rho}_n$) (fig. S3). Although not as great as the difference in $\hat{\rho}_n$ values for the YLS and ALP subjects, the UCB subjects' $\hat{\rho}_n$ values are skewed to the right relative to ALP subjects, indicating a greater efficiency orientation. The median $\hat{\rho}_n$ values are 0.259 and 0.005 for UCB and ALP subjects, respectively; a Wilcoxon rank-sum test rejects equality of the distributions ($P < 0.001$) (25).

Concluding remarks

People from all walks of life implement their distributional preferences in the real world. This is especially true for the elite YLS students in our sample, many of whom will assume positions of substantial power in economic and political affairs. We decomposed distributional preferences into two qualitatively different components—fair-mindedness and equality-efficiency tradeoffs—and measured both at the individual level in diverse samples of varying degrees of eliteness. Our experiment enabled us to distinguish fair-mindedness from equality-efficiency tradeoffs in the laboratory and to assess the extent of efficiency orientation in subject pools with different degrees of eliteness.

The increase in wealth and income inequality within and across countries is one of the defining social, economic, and political challenges of our time. Our results offer a potential new explanation for the muted policy response to increased income inequality in the United States: The equality-efficiency tradeoffs of the policy-making elite are such that they are far less inclined than is the general population to sacrifice efficiency to promote equality. As Gilens and Page (26) found, the preferences of the economic elites are far more correlated with public policy choices than are the preferences of the general public. Although there are many factors that contribute to the limited distributional response to rising inequality in the United States (ranging from loss-aversion, to attitudes toward fair-treatment of oneself by others, to moral hazard concerns, to beliefs about the extent of inequality), we focus on one potential cause: By favoring efficiency over equality, policy-makers may be acting on their own distributional preferences, which may be closely aligned with the interests and preferences of other members of the elite. The connections we draw between laboratory results and the degrees of eliteness promise to help in understanding the policies and practices that are implemented by the elite or the establishment in the broader world, and the experimental techniques and results that we have already developed provide promising tools for future work in this area.

Our results contribute to the broader discussion of the interplay between distributional preferences and tax policy. The vast and growing body of work on this topic includes theoretical analyses, experiments in the laboratory

and in the field, and surveys of distributional preferences, as well as the related but distinct notion of preferences for government redistribution. A recent study emphasizes the critical role that distributional preferences should play as a determinant of distributive policies generally and optimal tax policy in particular (27). This entire body of work overlooks a critical mediating factor between the measured distributional preferences of the general population and implemented policy choices: namely, the extent to which the distributional preferences of those in power differ from the preferences of voters. Ours is the first study to emphasize the existence of such differences, laying the groundwork for better understanding why public policy outcomes may diverge from stated voter preferences. We thus open a new and important window onto familiar mysteries, including the modesty of the policy response to rising inequality even in the face of growing popular outrage (28).

REFERENCES AND NOTES

- The term "affluenza" has entered the English language to describe the phenomenon of growing elite entitlement and excessive consumption, and academic work has connected higher socioeconomic class to a greater expressed sense of entitlement and greater narcissism (29). We conceptualize elite status in terms of educational attainment and lifetime wealth.
- All nine sitting Supreme Court justices and two of the past three presidents (as well as a frontrunner to become the next president) are graduates of either Yale or Harvard Law Schools. Most YLS students are also destined for membership in the economic elite, creating another potential channel of policy influence.
- A. M. Okun, *Equality and Efficiency: The Big Tradeoff* (Brookings Institution Press, Washington, DC, 1975).
- J. Andreoni, J. Miller, Giving according to GARP: An experimental test of the consistency of preferences for altruism. *Econometrica* **70**, 737–753 (2002). doi: [10.1111/1468-0262.00302](https://doi.org/10.1111/1468-0262.00302)
- R. Fisman, S. Kariv, D. Markovits, Individual preferences for giving. *Am. Econ. Rev.* **97**, 1858–1876 (2007). doi: [10.1257/aer.97.5.1858](https://doi.org/10.1257/aer.97.5.1858)
- R. Fisman, P. Jakiela, S. Kariv, The distributional preferences of Americans. *National Bureau of Economic Research* **20145** (2014); available at www.nber.org/papers/w20145.
- Detailed demographic data are not available for the YLS student body as a whole, but the evidence that we do have indicates that our YLS sample is broadly representative of the student body. Of YLS students graduating between 2007 and 2015, 46.7% were women, and 46.6% of our YLS subjects are female. Between 2007 and 2015, the average age at graduation was 28, whereas the average age in our YLS sample is 25.4 (but subjects are drawn from all 3 years of law school, and 1st-year students are slightly overrepresented).
- In our ALP elites versus nonelites analysis, we omit 157 of the 309 ALP subjects aged 40 and under from the experiment of (6) who did not report their household incomes, provided inconsistent responses to the two-part income question in the ALP survey, were unemployed, or did not meet our criteria for (either) eliteness or noneliteness (for example, because they held graduate degrees but reported low household incomes).
- R. Fisman, P. Jakiela, S. Kariv, How did distributional preferences change during the great recession? *J. Public Econ.* **128**, 84–95 (2015). doi: [10.1016/j.jpubeco.2015.06.001](https://doi.org/10.1016/j.jpubeco.2015.06.001)
- Materials and methods are available as supplementary materials on Science Online.
- S. N. Afriat, The construction of utility functions from expenditure data. *Int. Econ. Rev.* **8**, 67–77 (1967). doi: [10.2307/2525382](https://doi.org/10.2307/2525382)
- H. R. Varian, The nonparametric approach to demand analysis. *Econometrica* **50**, 945–972 (1982). doi: [10.2307/1912771](https://doi.org/10.2307/1912771)
- H. R. Varian, Non-parametric tests of consumer behavior. *Rev. Econ. Stud.* **50**, 99–110 (1983). doi: [10.2307/2296957](https://doi.org/10.2307/2296957)
- Let $\{(p^i, \pi^i)\}_{i=1}^{50}$ be the data generated by some individual's choices, where p^i denotes the i th observation of the price vector, and π^i denotes the associated allocation. An allocation π^i is directly revealed preferred to π^j , denoted $\pi^i R^D \pi^j$, if $p^i \cdot \pi^i \geq p^j \cdot \pi^j$. An allocation π^i is revealed preferred to π^j , denoted $\pi^i R \pi^j$, if there exists a sequence of allocations $\{\pi^k\}_{k=1}^K$ with $\pi^1 = \pi^i$ and $\pi^K = \pi^j$, so that $\pi^k R^D \pi^{k+1}$ for every $k = 1, \dots, K - 1$. In this notation, GARP requires that if $\pi^i R \pi^j$, then $p^j \cdot \pi^j \leq p^i \cdot \pi^i$ (that is, if π^i is revealed preferred to π^j , then π^j must cost at least as much as π^i at the prices prevailing when π^i is chosen).
- Choi *et al.* (30) provide more details on the power of the GARP test. Bronars (31) builds on Becker (32) to provide a test based on a comparison of the behavior of actual subjects to the behavior of simulated subjects who randomize uniformly on each budget line. The power of Bronars' test is defined to be the probability that a random subject violates GARP. In our experimental design, all random subjects had violations, implying that the Bronars criterion attains its maximum value.
- S. N. Afriat, Efficiency estimates of production functions. *Int. Econ. Rev.* **13**, 568–598 (1972). doi: [10.2307/2525845](https://doi.org/10.2307/2525845)
- The multiple decisions that subjects make could lead to their conceiving of their choices as a portfolio. As a result, equality-focused subjects might make choices that are less equality-focused by allocating more than their conception of fairness requires to self if the relative price of redistribution exceeds one and more to other if the relative price of redistribution is less than one, evaluating their portfolio of choices according to its expected payoff. Because a single decision round is used for payoffs, an equality-focused subject that follows this approach will create extreme inequality ex post. This approach also is not supported by the experimental data; very few subjects allocate all of the tokens to self when the relative price of redistribution is greater than one, or none of the tokens to self when the relative price of redistribution is less than one. More generally, it is possible that a subject implementing a preference over the rounds approached as a portfolio might appear, on account of our focusing on separate rounds, to be implementing a different preference. But only unusual preference orderings have the property that when implemented over portfolios, they appear as the consistent implementation of different preferences round by round. Many preferences over portfolios will instead generate inconsistent choices analyzed round by round. Our subjects' almost uniformly high CCEI scores thus constitute strong evidence against the hypothesis that they approach the rounds as a portfolio.
- Following Bronars (31), we compared the behavior of our actual subjects with the behavior of simulated subjects who randomize uniformly on each budget line. The mean CCEI in a sample of 25,000 simulated subjects is only 0.60; the vast majority of YLS subjects and many ALP subjects have CCEIs above 0.90, whereas none of the simulated subjects had a CCEI that high.
- We found similar results when we screened out selfish subjects who allocate themselves an average of more than 95% of the tokens.
- We found nearly identical results when different $\hat{\alpha}_n$ thresholds were used to classify subjects as fair-minded and selfish. Because the CES form approaches the Cobb-Douglas function $\pi_n^\alpha \pi_o^{1-\alpha}$ as $\rho \rightarrow 0$ (so that the expenditure shares are constant for any price of redistribution p), we only consider $\hat{p}_n = 0$ as the threshold between preferences weighted toward efficiency versus equality.
- J. G. Altonji, T. E. Elder, C. R. Taber, Selection on observed and unobserved variables: Assessing the effectiveness of Catholic schools. *J. Polit. Econ.* **113**, 151–184 (2005). doi: [10.1086/426036](https://doi.org/10.1086/426036)
- Of the two with no career data, one had, according to LinkedIn, left his or her last place of employment in 2011 but provided no further information. For the second subject, we could find no information whatsoever.
- We did not perform a parallel analysis for the parameter measuring fair-mindedness. Fair-mindedness may implicate payoffs beyond monetary income. Although elite jobs in the nonprofit, government, and academic sectors pay less than elite corporate sector jobs pay, they may offer substantial compensating personal benefits in the form of reputation. We therefore do not think it sensible to associate corporate employment with selfishness and nonprofit employment with fair-mindedness.
- In Table S1, we replicate the regression analyses reported in Table 2 with the ALP elite versus nonelite samples. We did not find a statistically significant association between ALP elite and equality-efficiency tradeoffs at the 50th percentile ($P = 0.224$), although we found a statistically significant association at the 75th percentile ($P = 0.016$). The differences between elites and nonelites are also insignificant when we look at CCEI scores and estimated $\hat{\alpha}_n$ parameters, but the coefficients are directionally consistent with higher consistency and lower fair-mindedness among elite subjects. Adding controls for gender and age has little effect on the coefficient estimates or significance levels.
- In Table S2, we replicate the regression analyses reported in Table 2 using the UCB and ALP samples. The differences between the UCB and ALP subjects are always in the same direction as the differences between the YLS and ALP subjects and are always statistically significant. We cannot include controls because we were unable to obtain individual-level data on UCB subjects owing to privacy concerns.
- M. Gilens, B. I. Page, Testing theories of American politics: Elites, interest groups, and average citizens. *Perspect. Polit.* **12**, 564–581 (2014). doi: [10.1017/S1537592714001595](https://doi.org/10.1017/S1537592714001595)
- E. Saez, S. Stancheva, Generalized social marginal welfare weights for optimal tax theory. *National Bureau of Economic Research* **18835** (2013); available at www.nber.org/papers/w18835.
- I. Kuziemko, M. I. Norton, E. Saez, S. Stancheva, How elastic are preferences for redistribution? Evidence from randomized survey experiments. *National Bureau of Economic Research* **18865** (2013); available at www.nber.org/papers/w18865.
- P. K. Piff, Wealth and the inflated self: Class, entitlement, and narcissism. *Pers. Soc. Psychol. Bull.* **40**, 34–43 (2014). doi: [10.1177/0146167213501699](https://doi.org/10.1177/0146167213501699); pmid: [23963971](https://pubmed.ncbi.nlm.nih.gov/23963971/)
- S. Choi, R. Fisman, D. Gale, S. Kariv, Revealing preferences graphically: An old method gets a new toolkit. *Am. Econ. Rev.* **97**, 153–158 (2007). doi: [10.1257/aer.97.2.153](https://doi.org/10.1257/aer.97.2.153)
- S. G. Bronars, The power of nonparametric tests of preference maximization. *Econometrica* **55**, 693–698 (1987). doi: [10.2307/1913608](https://doi.org/10.2307/1913608)
- G. S. Becker, Irrational behavior and economic theory. *J. Polit. Econ.* **70**, 1–13 (1962). doi: [10.1086/258584](https://doi.org/10.1086/258584)

ACKNOWLEDGMENTS

We are grateful to J. Andreoni, D. Bernheim, S. DellaVigna, C. Jolls, D. Kahn, B. Polak, M. Rabin, E. Saez and A. Schwartz for helpful discussions and comments. This paper has also benefitted from suggestions by the participants of seminars at several universities and conferences. The experiments reported in this paper were funded by the Columbia University Graduate School of Business, the Center for Equitable Growth at UCB, YLS, and the Oscar M. Reubhausen Fund. S.K. acknowledges the National Science Foundation and the Multidisciplinary University Research Initiative (MURI) for financial support. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of any funding agency. Replication data files are available online at <https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi%3A10.7910%2FDFN%2FJYHMRA>.

SUPPLEMENTARY MATERIALS

www.sciencemag.org/content/349/6254/aab0096/suppl/DC1
Materials and Methods
Figs. S1 to S3
Tables S1 and S2
References (33–36)

16 March 2015; accepted 10 August 2015
10.1126/science.aab0096

This copy is for your personal, non-commercial use only.

If you wish to distribute this article to others, you can order high-quality copies for your colleagues, clients, or customers by [clicking here](#).

Permission to republish or repurpose articles or portions of articles can be obtained by following the guidelines [here](#).

The following resources related to this article are available online at www.sciencemag.org (this information is current as of September 25, 2015):

Updated information and services, including high-resolution figures, can be found in the online version of this article at:

<http://www.sciencemag.org/content/349/6254/aab0096.full.html>

Supporting Online Material can be found at:

<http://www.sciencemag.org/content/suppl/2015/09/16/349.6254.aab0096.DC1.html>

A list of selected additional articles on the Science Web sites **related to this article** can be found at:

<http://www.sciencemag.org/content/349/6254/aab0096.full.html#related>

This article **cites 17 articles**, 3 of which can be accessed free:

<http://www.sciencemag.org/content/349/6254/aab0096.full.html#ref-list-1>

This article appears in the following **subject collections**:

Economics

<http://www.sciencemag.org/cgi/collection/economics>