Distributional preferences in larger groups: Keeping up with the Joneses and keeping track of the tails

Raymond Fisman, Ilyana Kuziemko and Silvia Vannutelli*

April 22, 2020

Abstract

We study distributional preferences in larger "societies." We conduct experiments via Mechanical Turk, in which subjects choose between two income distributions, each with seven (or more) individuals, with hypothetical incomes that aim to approximate the actual distribution of income in the U.S. In contrast to prior work, our design allows us to flexibly capture the particular distributional concerns of subjects. Consistent with standard maximin (Rawlsian) preferences, subjects select distributions in which the bottom individual's income is higher (but show little regard for lower incomes above the bottom ranking). In contrast to standard models, however, we find that subjects select distributions that *lower* the top individual's income, but not other high incomes. Finally, we provide evidence of "locally competitive" preferences—in most experimental sessions, subjects select distributions that lower the income of the individual *directly* above them, while the income of the individual *two* positions above has little effect on subjects' decisions. Our findings suggest that theories of inequality aversion should be adapted to account for individuals' aversion to "topmost" and "local" disadvantageous inequality.

JEL Classification Numbers: C91, D63, H23.

Key words: Inequality aversion; Envy; Redistribution

*For helpful feedback we thank John Cisternino, David Moss, Michael Norton, and seminar participants at Yale. Financial support from the Tobin Project is gratefully acknowledged. Fisman: Boston University and NBER (email: rfisman@bu.edu); Kuziemko: Princeton University and NBER (email: kuziemko@princeton.edu); Vannutelli: Boston University (email: svann@bu.edu)

1 Introduction

Economists have long recognized that individuals incorporate others' payoffs into their own utility. This insight has given rise to a rich theoretical and experimental literature to better understand the structure of individuals' distributional preferences. Economists' interest in the topic is not merely academic — models of distributional preferences can help inform our understanding of support for redistributive policies, political preferences, and the provision of public goods. The relevance of understanding distributional preferences has risen to particular prominence, given the secular increase in inequality since the late 1970s (see, e.g., Piketty and Saez, 2003, Piketty et al., 2018).

Much of the emphasis has been on testing various models of inequality aversion. Many of these models assume functional forms in which the incomes of others matter only in aggregate (as in, for example, Bolton and Ockenfels (2000), which argues that individuals care about their own level of income and their *share* of the total), or which add the possibility that individuals try to help the worst-off person (as in Rawlsian preferences, explored in Charness and Rabin (2002) and Engelmann and Strobel (2004)). Fehr and Schmidt (1999) allow the disutility created by income gaps to depend on whether the gaps are disadvantageous (resulting from incomes above the individual) or advantageous (resulting from incomes below the individual), but within these two groupings income differences are just aggregated to total disadvantageous and advantageous inequality.

These models have performed well in experimental tests in which the decision-maker divides earnings between herself and a single recipient or pair of recipients. But relatively little work has explored their predictions in larger groups that may have more direct analogs in the real world, and when larger-group experiments have been run, they have generally imposed a strict formula on redistribution (most commonly, proportional taxation with lump sum redistribution). Not only are such forms of redistribution more restrictive than those observed in the real world, they also prevent subjects from expressing concerns for particular individuals or groups in society. For example, as evidenced by calls for redistributing wealth away from the top "one percent," there may be particular concern for disadvantageous inequality focused on society's highest earners.

To more fully characterize distributional preferences, we devise a simple experiment that allows subjects to express, via revealed preference, their concern for inequality at different points in the income distribution relative to their own. As a result, we may distinguish the extent to which subjects place equal weight on the income gaps between themselves and *all* other individuals above them (as in Fehr and Schmidt (1999)) or if they place higher weight on the incomes of people in certain positions (e.g., the very top). Similarly, we can test whether all advantageous income gaps between oneself and those below carry equal weight.

We conduct a set of experiments via Mechanical Turk (MTurk) in which each subject is confronted with a choice between two hypothetical societies A and B, each with a different income distribution. In most versions of the experiment, each distribution is comprised of seven individuals, including the subject herself. The two societies have different income distributions, generated by taking independent draws from the same underlying process. The data-generating processes are designed to be reflective of the rough level of prosperity in the United States, but to have some tilt toward the upper right tail, given distributional preferences over this group will have the largest tax implications as a result of their disproportionate share of income. For example, a typical distribution would be {\$10,934, \$28,102, \$62,275, \$92,479, \$107,973, \$151,869, \$188,371}). By construction, in most variants, the subject's *own* income is held constant in Societies A and B (e.g., if she were the fourth-ranked person in the given example, her income would be \$92,479 and have a rank of four in *both* distributions, whereas the other values would vary, independently sampled from the given generating process). We make this choice to emphasize the role of *others'* payoffs in choosing distributional outcomes. Our design allows us to distinguish, for example, whether individuals put more weight on reducing inequality at extreme income levels such as the top and bottom, or focus on inequality nearer to the subject's own income.

Our first set of results does not explicitly consider the position of the subject herself, but focuses instead on whether certain positions – most obviously the highest- and lowest-ranked ones – play a particularly prominent role in subjects' choices. We begin non-parametrically, showing the effect of every position's income value on the subject's decision between the two distributions. We find a very robust emphasis on reducing extreme inequality. Consistent with Rawlsian preferences and the results in Charness and Rabin (2002) for two- and three-subject settings, subjects are significantly more likely to select the distribution that raises the bottom individual's income. This effect is very large: a subject is about 30 percentage points more likely to select the distribution in which the least well-off individual's income is higher. More novel, we also find a robust and quantitatively important emphasis on *lowering* the income of the individual in the highest position in the distribution—subjects are more than ten percentage points more likely to select the distribution with a lower income in the top position, all else equal. Apart from the top and bottom incomes, no other absolute position has any impact on subjects' decisions.

In our second set of results, we define others' positions in a *relative* sense: one position above the subject, one below the subject, and so forth. We observe a large and significant desire to reduce the income of the individual in the position *directly above* the subject's own income. In fact, this effect is comparable in magnitude to subjects' preference for lowering the income of the individual at the highest position in the distribution. By contrast, the income *two* positions above her has no impact on the choice of distribution, and we can reject at high levels of precision that these two effects are of equal magnitude. Such a result is inconsistent with a general desire to reduce inequality.

To organize our findings, we introduce the distinction between *locally competitive* versus *topmost competitive* preferences to reflect our subjects' particular focus on incomes very close to their own as well as those at the top tail of the distribution. This framing can provide some (parsimonious) guidance on the weights that individuals place on inequalities at particular ranks in the income distribution, thereby enriching models of inequality aversion that are standard in the literature. In particular, we can reject that individuals treat disadvantageous income gaps symmetrically: instead we show that the immediate income gap between the subject and the person right above her as well as the income gap between her and the richest person matter more than all other disadvantageous gaps. Similarly, we can reject that all advantageous inequality is equal: instead, the advantageous inequality between the subject and the worst-off person appears to differentially reduce utility.It is further interesting to note that topmost and bottom-most inequality aversion in our data are stronger for Democrats than Republicans, while locally competitive preferences are invariant to political affiliations, indicating that something beyond ideology is driving the latter result.

Our framework may also help to reconcile some attitudes toward inequality that are harder to explain with standard models. For example, consideration of local versus topmost competitiveness is consistent with the popular outrage over the high incomes of the top one percent. It can similarly explain why people care about "keeping up with the Jones" while at the same time ignoring the somewhat more prosperous Johnsons.[1]

To summarize, while our non-parametric approach to examining social preferences in larger groups might have led to unwieldy results, in practice, the patterns we document are intuitive, straightforward to summarize, and can be expressed via parsimonious adjustments to existing models.

In the years following our initial pair of experiments (conducted in September, 2013), we ran a number of additional "sessions" that subjected our analysis to a wide range of robustness checks.[2] We allowed the subject's *own* income to differ across the two distribu-

[1]For example, Luttmer (2005) shows that within relatively small geographic units, average local income negatively predicts individual well-being, holding one's own income constant. As these units are more economically homogeneous than the entire country, the result suggests that individuals care about the incomes of those close to them in the distribution (though that interpretation is confounded by geographic proximity).

[2]Given we use MTurk and not a true lab, "sessions" is a slight abuse of the language, but by "session" we mean separate MTurk surveys administered at particulars dates and times.

tions, varied the generating processes to create distributions with higher levels of inequality, and increased the number of individuals in each "society" to nine. We also varied the way that the distributions were presented to subjects, pulling the bar representing the subject's own income away from those representing the incomes of others. Up to this point, all experiments confronted subjects with *hypothetical* choices between pairs of income distributions. In a real-stakes version, subjects were informed that, with ten percent probability, their choice would be implemented for stakes equal to one ten-thousandth of the income distributions presented, with randomly drawn MTurk workers as recipients. Thus, for example, if an individual in a selected distribution was assigned an income of \$140,000 and that distribution was selected for real payoffs, a random MTurk worker would receive \$14 as payment. Finally, in 2019 we implemented three further variants, which served to further assess the sensitivity of our results to framing concerns: a "democracy" treatment in which subjects' decisions were presented as a vote on their preferred society, rather than a decisive choice as in a dictator game; a "histogram" treatment in which each society was presented as a histogram over a very large number of individuals, in order to capture distributional concerns in a setting that more closely reflects actual societies; and a "labeled mean" treatment in which a line capturing a society's mean income was provided for each graph.

We find that subjects in every session tended to choose distributions that reduce the income of the individual in the highest position (i.e., subjects always exhibit topmost competitiveness). This effect is largest for the very-high-inequality variant of the experiment, suggesting that the focus on the top is greatest when inequality is higher and not merely an artifact of the highest position catching subjects' attentions. Subjects also increase the income of the individual in the lowest position. Evidence for "locally competitive" preferences is also observed in every variant apart from the "real stakes" and "labeled mean" sessions. We return to discuss some possible explanations for these exceptions when we present our experimental findings.

Our paper contributes most directly to the recent literature that aims to infer redistributive preferences via survey methods. These studies have generally been devised to better understand attitudes toward (and impediments to) inequality-reducing redistribution in general (e.g., Kuziemko et al. (2015), Norton and Ariely (2011)), without exploring the particular structure of these preferences.

More broadly, our work builds on the large body of research that aims to characterize the nature of distributional preferences. In our experimental design, we attempt to bridge social preferences as typically studied in small groups with small stakes to preferences over more policy-relevant income distributions. The literature we build upon encompasses the theoretical contributions referenced earlier, many of which have experimental components to them involving just one or two recipients.[3]

Our paper joins a smaller literature on distributional preference experiments involving large groups. As noted earlier, these studies tend to impose a formulaic redistribution parameter such that all poorer individuals are made better off and all richer individuals made worse off, potentially subject to some efficiency loss (see, in particular, Durante et al. (2014), Ackert et al. (2007), and Beckman et al. (2004)). In contrast to our work, these prior studies do not separate subjects' concerns for others' incomes at particular points in the distribution. Importantly, Durante et al. (2014) focuses on redistribution via a flat tax, which is both rare in the real world, and also cannot pick up aversion to very high incomes, which we find to be important in practice.[4]

As our subjects are in a sense acting as social planners (though they are "planning" a society of which they are a member), our work relates to the optimal tax literature under non-standard preferences (either of the social planner herself or of individuals in society). Recently, several important contributions on the theory side of this question have emerged. Saez and Stantcheva (2016) examines optimal tax outcomes when the social planner is loss-averse, Farhi and Gabaix (2015) develops optimal tax formulae under a number of behavioral anomalies (e.g., inattention and mental accounting), and Lockwood (2016) focuses on present bias. Our approach, as in Charité et al. (2015), who test whether social planners respect individuals' reference points, is more experimental.[5]

[3]See Kahneman et al. (1986) for the earliest dictator experiment and Forsythe et al. (1994) for the first appearance of the standard dictator game. More recent research has explored how dictators' fairness principles are affected by considerations such as deservingness (e.g., Almås et al. (2010), Krawczyk (2010)) and extrinsic versus intrinsic motivations (Cappelen et al. (2017)). Recent research has also generalized the dictator framework to incorporate different prices of giving to create a tradeoff between equality and efficiency (see Andreoni and Miller (2002) and Fisman et al. (2007)). We see our work as extending this tradition to consider the more complex set of tradeoffs that come with distributional choices involving multiple others.

[4]While less directly related to distributive preferences, a recent paper examines decision-making in groups far larger than those in typical experiments. Schumacher et al. (forthcoming) finds that many subjects make welfare-decreasing decisions while acting as social planners for large (up to 32 individuals) groups in lab experiments, because subjects weigh a salient benefit for a minority of the group as more important than a small cost for the majority of the group, even if the sum of the costs is greater than the sum of the benefits.

[5]In comparing our results from hypothetical and real-stakes settings, we also contribute to a small literature that seeks to test whether non-incentivized results generalize to incentivized settings. In general, earlier research on this topic has found mixed results. Camerer and Hogarth (1999) perform a meta-analysis of 74 studies that either have no, low or high-powered incentives. They find that the effect of real stakes depends on the experimental task. Beattie and Loomes (1997) compare three payment schemes: hypothetical, randomly picking one of several questions for payment, and paying out for each question. They find that choices involving *pair-wise* comparisons of lotteries are not affected by payment (although subjects are less likely to violate expected utility theory over

The remainder of the paper is organized as follows. Section 2 develops a very simple generalization of the Fehr-Schmidt model that we will use to guide our empirical analysis. Section 3 describes the experimental design. Section 4 explains the data collection procedure. Section 5 presents the results from our main experiment, followed by Section 6, which reports on the results from our follow-up sessions. Section 7 offers concluding thoughts and suggestions for future work.

2 Standard models of distributional preferences

In this section, we summarize the classic inequality aversion model of Fehr and Schmidt (1999), and also present a more flexible (though less parsimonious) version that will serve as a point of departure for our empirical exercise.

We begin with the Fehr-Schmidt model. Consider a society of n individuals and a vector of payoffs $\mathbf{x} = (x_1, ..., x_n)$; the utility function of individual i is then given by:

$$U_i(\mathbf{x}) = x_i - \frac{\alpha_i}{n-1} \sum \max\{x_j - x_i, 0\} - \frac{\beta_i}{n-1} \sum \max\{x_i - x_j, 0\}, \tag{1}$$

where α_i is individual i's aversion to disadvantageous inequality and β_i is his aversion to advantageous inequality. The model assumes $0 \leq \beta_i < 1$ and $\alpha_i \geq \beta_i$. These conditions imply that individuals do not like advantageous or disadvantageous inequality ($\alpha_i,\ \beta_i \geq 0$), but that they are not willing to reduce own-income in order to reduce inequality ($\beta_i < 1$), holding others' incomes constant. The second assumption implies additionally that individuals dislike falling behind more than they dislike being ahead of others.

In our setting, in which own-income is held constant (though we relax this constraint in one experimental session), this model predicts that subjects will choose the distribution that minimizes the aggregate payoff differences relative to others, with a greater weight on decreasing high incomes rather than increasing low incomes. However, the model makes no distinction among all individuals above the subject's income or all individuals below: all incomes within each group are given the same weight, $\frac{\alpha_i}{n-1}$ or $\frac{\beta_i}{n-1}$, respectively.[6]

Individual- or rank-specific comparisons can easily be embedded in the Fehr-Schmidt model by allowing parameters to vary for each of the other individuals' positions. Without

complicated sequences of lotteries when they are paid for *each* question). By contrast, while our pair-wise comparisons are not tests of rationality, we do find that payment matters for subjects' decisions. Etchart-Vincent and l'Haridon (2011) find that hypothetical and incentivized choices do not differ for the choice to bear risk in the loss domain, but that hypothetical choices in the gain domain are more risk-seeking than incentivized choices (consistent with Holt and Laury (2002)).

[6]Bolton and Ockenfels (2000) adopt an approach that similarly yields no prediction of either locally or topmost competitive preferences, based on the utility function $U_i(x) = U(x_i, \frac{x_i}{\sum x_j})$

loss of generality, we order the vector of incomes $\mathbf{x}$ in increasing order, so that the fully flexible Fehr-Schmidt model may be expressed as:

$$U_i(\mathbf{x}) = x_i - \frac{1}{n-1}\sum_{j=i+1}^{n} \alpha_{i,j}(x_j - x_i) - \frac{1}{n-1}\sum_{j=1}^{i-1} \beta_{i,j}(x_i - x_j) \tag{2}$$

Our experimental design allows us to potentially estimate all of these individuals' specific weights ($\alpha_{i,j}$ and $\beta_{i,j}$). If we allowed for a fully flexible specification, it would lead to a very large number of parameters (since they potentially depend on where the subject is in the income distribution). As a result, we will present a set of graphs at the beginning of Section 5 that displays each of these parameters. We then inspect the patterns to determine the regression specifications that the data appear to suggest. While there is some subjectivity in going from inspection of the graphs to an explicit regression specification, we will be aided in this endeavor by the fact that we observe very clear patterns in the data: subjects' inequality aversion focuses on the top and bottom individuals' incomes (captured by $\alpha_{i,n}$ and $\beta_{i,1}$ respectively), and locally disadvantageous inequality aversion (captured by $\alpha_{i,i+1}$).

3 Experimental Design

The centerpiece of the survey presents each subject with a binary choice between two income distributions, which are called "Society A" and "Society B." The survey experiment begins with the following initial instructions (or a close variant of them, depending on whether we were running the main survey experiment or one of the modified versions that we ran to explore the robustness of our various findings):

> In each round you will see two graphs displayed on your screen. Each graph represents a distribution of payoffs that you can choose to assign to yourself and to the other participants in your group. In each round, you must decide which distribution, Society A or Society B, you prefer.

Subjects then completed a practice round, which was accompanied by the following instructions:

> Two graphs are displayed below. Each graph represents a distribution of payoffs that you can choose to assign to yourself and the other participants in your group. The red bar in each graph indicates your position and payoff in the group. Please select which distribution you prefer.

After completing the practice round, subjects confirmed that they had read and understood the directions before completing the ten subsequent iterations that constitute the data we use in our analysis (see the screenshot in Figure 1).[7]

In every iteration, the subject's own position in the income distribution was selected at random. Further, in each iteration, instructions were reprinted above the two graphs, as shown in the screenshot. Following the last iteration, subjects completed a short survey on their attitudes toward government redistribution, their political preferences and voting decisions, and basic demographics like age, gender, and income.

We focus our presentation of the results from the initial pair of experimental sessions that we conducted on September 9$^{\text{th}}$ and September 17$^{\text{th}}$, 2013. In both sessions, subjects were presented with choices that took the precise form illustrated in Figure 1, differing only in the process by which the income distribution values were generated.

The income values for these sessions were drawn from uniform distributions in each of seven ranges: $(a_1, b_1), (a_2, b_2)...(a_7, b_7)$, where the subscripts denote the position $p \in \{1, 2, ..., 7\}$ in the distribution. Note that *position*, as we define it, is *increasing* in income, as opposed to *rank*. We set $b_p = a_{p+1}$, so that the union of the intervals is $(a_1, b_7) \setminus \{b_1, b_2, ..., b_6\}$ (i.e., the full interval minus a subset of measure zero). The non-overlapping ranges ensure that in no case can ranks change from Society A to Society B. As such, if the subject finds herself in position 4 in Society A, she will be in position 4 in Society B. We made this choice to simplify the setting: the person right above a subject may move closer or further away, but the subject can never "leapfrog" over him (nor can the subject be leapfrogged by the person directly below).

To probe robustness, in our main experimental sessions we vary how the a_p and b_p values are set. In one case, which we term "Absolute Differences" (AD), the ranges were kept constant, in \$20,000 increments, beginning at \$10,000 (i.e., 10,000–30,000, 30,000–50,000. . . 130,000–150,000). To give some sense of where these values sit in the distribution of U.S. pre-tax, pre-transfer income, the midpoint of the lowest interval is at roughly the 24th percentile and the midpoint of the highest interval is at roughly the 94th.[8] In the second case, which we term "Percentage Differences" (PD), we keep the percentage increase in a comparable range across positions in the income distribution, increasing the range by \$5,000 at each level (i.e., the ranges are 10,000–25,000, 25,000–45,000. . . 175,000–220,000). As indicated in the instructions, the subject's own income is presented in a different color in

[7] Subjects were additionally assured that all responses would remain anonymous.

[8] These percentiles are based on the 2016 CPS. The midpoint of the second interval is at roughly the 52nd, that of the third at the 72nd, that of the fourth at the 82nd, that of the fifth at the 88th, and that of the sixth at the 92nd. As noted in the introduction, we wanted distributions with some skew to the right.

both distributions, and in all cases the subject's own income is identical in each distribution to focus the decision-maker on inequality rather than own-income.

We conducted a number of variants on this basic design to probe the robustness of our results to different income distributions and ways of presenting them to subjects. These variants allow *own* income to vary; provide alternative presentations of the income distributions to assess the extent to which the results are driven by the particular manner in which income distributions are presented; change the distribution from which the income values are drawn; make the experiment for "real money;" and allow subjects to vote for their preferred society rather than choosing directly.

Appendix Table A.1 provides a full list of the treatments (the main experiments plus the companion experiments) as well as the dates they were conducted.

The interested reader can take the full experiment online at `nautech-clients.com/tobin/survey12/`. The version posted online is the "real stakes" session (described in detail in Section 6; its instructions are virtually identical to those of the main experiment, with the addition of a screen which explains how payoffs will be determined as a result of the subject's choices in the experiment).

4 Data

4.1 Data collection

In recent years, social scientists have increasingly used MTurk to perform experiments and collect survey data (see Kuziemko et al., 2015 and papers cited therein for a review). We briefly describe our data-collection procedures in this subsection, but relegate details to Appendix D, including a discussion of the steps we took to ensure the quality of our data.

To limit selection bias while also giving workers an honest description of the task, we provided a short, neutral description of our survey ("This survey is part of an academic research survey") that could be viewed by MTurk workers before they signed up to participate. To limit heterogeneity of the sample, we collected all data on workdays during daylight hours on the East Coast of the United States. Compensation was set to $.50 which, given the median completion time of seven minutes, works out to an hourly wage of $4.25. Though we cannot find official data on average wages on MTurk, reading through worker forums suggests that we are paying a generous wage (and indeed our posted surveys were always filled within a short period of time).

We collected data from ten separate sessions.[9]The sessions differ in the way that the

[9]While it would have been preferable to run all treatments concurrently, randomizing subjects into each arm, in practice we ran variants sequentially in response to feedback we received. The

income distributions were generated or presented, as detailed in the preceding section. We drop any subject who participated in a previous session, though we show in Appendix D that results are robust to including them.

Appendix D describes the various steps we took to ensure that our respondents were actual humans, living in the United States, taking the survey in good faith. Since 2018, researchers have found an increase in "bots" (algorithms that masquerade as human MTurk workers) on MTurk, so these steps are especially important for the more recent sessions. Basic cross-tabs of the data are reassuring (for example, subjects who report Republican party affiliation are roughly fifty percent richer than those who report Democratic affiliation and are more likely to be male).

We informed subjects upfront that the survey was part of an academic study. Given academia's left-wing reputation, one might worry that social-desirability bias would lead subjects to give more pro-redistribution answers (see, e.g., Bernardi, 2006, Dalton and Ortegren, 2011). In our setting, such concerns may be limited, as earlier research suggests that web-based surveys may be less prone to social desirability bias than traditional in-person interviews (Kreuter et al., 2008). We further tried to mitigate any such concerns by emphasizing early in the survey instructions that we sought individuals' genuine responses, explaining that: "You are invited to participate in an opinion survey. *There are no right or wrong answers* [emph. in original]. In one session, we also directly ask respondents if they found the survey biased, and the vast majority report they did not."[10]

4.2 Data sample

Table 1 provides details on all MTurk workers from all sessions that we conducted, comparing them to the (weighted) population of adults sampled in the 2014 General Social Survey. Consistent with past work using MTurk, we find that younger, male, and college-educated subjects are over-represented in our sample.[11] Also consistent with prior work, household incomes are relatively low among MTurk workers, especially considering their higher-than-average levels of education. On the social and political variables that may directly relate

consistent patterns we observe across the various treatments – conducted with substantial time lapses between them – alleviates this concern somewhat.

[10]In the "real stakes" variant, we asked respondents directly about whether they perceived some left-wing, right-wing or other sort of bias in our survey. Roughly 7.6 and 2.7 percent, respectively, said it was biased in a "politically liberal" or "politically conservative" manner, another 2.7 percent said it was biased in some other manner, with the remaining 87 percent saying they did not detect any bias. While these cross-tabs cannot speak to whether subjects were biased in some subconscious manner, we are somewhat reassured that social desirability bias is unlikely to be large.

[11]Unfortunately our post-experiment survey did not ask for respondents' race, but other studies have found that MTurk workers are less likely to be minorities than the U.S. average.

to distributional preferences, our sample is more likely to have voted than the average GSS respondent. The two samples are nearly identical in their opinion of how much the government should reduce income differences through redistribution. Interestingly, however, MTurk subjects are more likely to believe that success is a matter of luck than hard work, relative to GSS respondents. In a robustness check we present results reweighted to be reflective of the GSS population based on age, gender, income, and belief that the government should reduce income differences.

4.3 Notation and definitions

Each observation in our data is a subject-iteration, with each subject confronting ten iterations. Each iteration consists of a choice between two distributions, and thus each observation has associated with it a number of comparisons between these two distributions. Before proceeding to our main specifications and results, it is useful to provide some notation and define several terms to facilitate our exposition in the next section.

We define several variables that capture differences between the two income distributions. Let $Income_p^D$ be the income of the individual in position $p = 1, \ldots 7$ in Society $D \in \{A, B\}$. Recall that *position* is increasing in income, so the poorest person in a seven-person distribution has position 1 and the richest person has position 7.

We define $DiffIncome_p$ as the income for position p in Society B minus the income for position p in Society A. The preferences described in Section 2 predict that subjects will make decisions based on *absolute* difference in income for p in B versus A, weighted by the importance that i attaches to each position p. Thus, Society B is preferred to Society A if and only if:

$$\sum_{j=p_i+1}^{7} \alpha_{j,p_i} DiffIncome_j - \sum_{j=1}^{p_i-1} \beta_{j,p_i} DiffIncome_j > 0 \tag{3}$$

We also employ an alternative measure of differences in income inequality between the two distributions that is not sensitive to the widely differing income ranges at different positions in the distribution. Specifically, instead of $DiffIncome$, we look at a binary indicator variable that captures simply whether the income for position p is higher in Society B than in Society A. That is,

$$SignIncome_p = \begin{cases} 1 & \text{if } Income_p^B - Income_p^A > 0 \\ 0 & \text{otherwise.} \end{cases}$$

Note that for expositional parsimony we engage in some abuse of notation by calling this variable *SignIncome*, when in fact it takes values of 0 and 1, not -1 and 1.

Finally, given our interest in testing whether respondents focus on those closer to themselves in the distribution, we also define measures that are *relative* to the subject's own position. We thus define, for subject in position p, $DiffIncome^{+1} = DiffIncome_{p+1}$. So, in the preceding example (illustrated in Figure 1), $DiffIncome^{+1} = 18,099$. [12]

Past work has found that subjects often try to maximize total surplus, so it is natural to consider total income as a control in some specifications (though its introduction into the Fehr-Schmidt model is not without complications—we return to this point below). We define $DiffSurplus$ as the difference in total income of all individuals in Society B versus Society A:

$$DiffSurplus = \sum_{r=1}^{7} Income_r^B - \sum_{r=1}^{7} Income_r^A.$$

Similarly, we generate an indicator variable, *SignSurplus*, that denotes whether Society B has greater aggregate income than Society A.

5 Main Results

We first present visual displays of the data to depict how subjects decide between the two distributions, and then proceed to more formal regression results. Specifically, in the next section we provide an initial set of results that mirror the fully flexible specification in equation (2), which we will use in large part to motivate the more parsimonious specification that we deploy in our main regression tables.

5.1 Graphical evidence

We begin by exploring how subjects' decisions depend on income differences between the two distributions, independent of the subject's own position. In Figure 3 we show the results from the following specification:

$$Choose_{ik}^B = \alpha + \sum_{q=1}^{7} \lambda_q SignIncome_{q,ik} + \boldsymbol{P_{ik}} + \epsilon_{ik}, \tag{4}$$

which includes seven fixed effects for the position held by i in decision k ($\boldsymbol{P_{ik}}$). $Choose_{ik}^B$ is an indicator variable for subject i in iteration k of the experiment choosing Society B, and

[12] We similarly define $DiffIncome^{+2} = DiffIncome_{p+2}$, $DiffIncome^{-1} = DiffIncome_{p-1}$, and $DiffIncome^{-2} = DiffIncome_{p-2}$.

$SignIncome_{q,ik}$ is an indicator variable for position q having a higher income value in Society B. Each coefficient λ_q can be interpreted as the percentage point increase in likelihood that the subject selects Society B if the income of position q is higher in B. We also graph 95 percent confidence intervals, using standard errors clustered by subject.

In general, inequality aversion will lead subjects to pick distributions in which low positions have relatively high incomes. The graph clearly indicates a concern for raising the income of the poorest member of society: the probability of selecting Society B is nearly 30 percentage points higher if the income in its lowest position is higher than in Society A. We also observe an important role for the highest income—subjects are more than 10 percentage points more likely to select Society B if its richest individual has a *lower* income. For positions two through six, we observe precisely estimated zero coefficients indicating that, on average, incomes in these positions had no effect on subjects' decisions. Overall, our findings indicate that models of inequality aversion may wish to account for extremes in income—both rich and poor—and place less emphasis on intermediate incomes.

We next explore whether subjects' choices are affected by incomes *relative* to their own, as would be the case in the Fehr-Schmidt model. We do so by allowing the coefficients in the preceding analysis to vary depending on the subject's own position in the distribution, so that for each $p \in \{1, 2, \ldots, 7\}$, we estimate the following equation via OLS:

$$Choice^B_{ik} = \alpha + \sum_{q \neq p} \eta^p_q SignIncome_{q,ik} + \epsilon_{ik} \tag{5}$$

The estimation in each case is for all decisions k made by subject i in which she was assigned position p in the income distribution. Similar to the preceding figure, the η^p_q coefficients tell us whether subjects are more or less likely to choose a distribution that is favorable to position q when subjects are themselves in position p.

We plot the estimated η^p_q coefficients separately for each value of p, across the seven panels of Figure 4. As expected given the patterns in Figure 3, regardless of assigned rank, for all $p > 1$, η^p_1 is large and positive, indicating that subjects in all positions put considerable weight on raising the income of the least well off individual (recall, η^p_1 is not defined for $p = 1$). We similarly observe that for all $p < 7$, η^p_7 is negative across all panels, indicating a general desire to "soak the rich."

The only other case for which we observe a significant deviation from zero across all panels is for the position directly above the subject's own. In every panel, the "one above" coefficient is negative and significantly different from zero at the five-percent level. No other coefficient in positions two through six is significant across all panels, regardless of its position relative to the subject. To emphasize the importance that subjects place on "one above" incomes in particular, in panels (a) - (d) we can compare concern for the incomes of those one and two positions above the subject's own. In each case, we observe that for each own-position p, $\eta^p_{p+1} < \eta^p_{p+2}$ (significant at least at the ten-percent level in all cases). That is, subjects are averse to picking the distribution in which the individual in position $p+1$ has a relatively high income, whereas the incomes of individuals in position $p+2$ are relatively unimportant. (In panel (e), we observe that $\eta^p_{p+1} > \eta^p_{p+2}$, but this comparison conflates the effects of topmost and local competitiveness.)

Other than these patterns, which indicate aversion to inequality at the extremes as well as local competitiveness, relative incomes in other positions are uncorrelated with subjects' choices. This pattern is difficult to reconcile with standard models of distributional preferences that emphasize aggregate differences or raising only extremely low incomes.

5.2 Regression results

Motivated by the preceding results, we present our main regression estimates in the following parsimonious specification.

$$\begin{aligned} Choose^B_{ik} = \beta_1 DiffIncome^{+1}_{ik} + \beta_2 DiffIncome^{+2}_{ik} + \beta_3 DiffIncome^{1}_{ik} \\ + \beta_4 DiffIncome^{7}_{ik} + \lambda X_{ik} + e_{ik}, \end{aligned} \tag{6}$$

where X_{ik} are covariates related to subject i or iteration k (e.g., subject fixed effects, iteration fixed effects), which we vary to probe robustness. This specification focuses our analysis on the patterns that emerged in the previous section, allowing us to explore the robustness of inequality aversion toward top, bottom, and "one above" incomes across a range of specifications (we include $DiffIncome^{+2}$ to ensure that, in looking at "just above" incomes, we distinguish local competition from general aversion to disadvantageous inequality). For these analyses, we pool all decisions in which subjects held positions two through five, so that all covariates are defined. Recall that we pool the first two experimental sessions, which constitute the "baseline" experiment before we explore variants of the experiment. Throughout, monetary values are expressed in units of $\$10,000$ to make the output tables more readable. The coefficient on $DiffIncome^1$, for example, may be interpreted as the percentage point increase in the probability of selecting Society B if the income of the poorest individual in Society B increases by $\$10,000$ relative to the income of the poorest individual in Society A.

We present the results from this specification in Table 2. In column (1) we show the results including the set of $DiffIncome$ variables. The coefficients on $DiffIncome^1$ and $DiffIncome^7$ are positive and negative, respectively, and both highly significant ($p < 0.0001$ in both cases). The coefficient on $DiffIncome^1$ is 0.195, implying that a one standard deviation increase in $DiffIncome^1$ (0.688) leads to a 13.4 percentage point greater probability that a subject selects Society B. The coefficient on $DiffIncome^7$, -0.0426, implies that a one standard deviation increase in $DiffIncome^7$ (1.231) leads to a 5.2 percentage point lower probability that a subject selects Society B. We also find a strong local competition effect: we estimate that $\beta_1 = -0.0371$, whereas $\beta_2 = -0.000224$. The difference is significant at the one percent level.

In column (2) we include 10 question-order (iteration) fixed effects, which has little impact on our estimates of the coefficients on the $DiffIncome$ variables. Column (3) includes fixed effects for the subject's position in the income distribution. Column (4) excludes subjects who completed the experiment very rapidly (less than 4 minutes). In all cases, the coefficients on the $DiffIncome$ variables are virtually unchanged.[13]

In Appendix Table A.2 we present results that reweight observations to be reflective of the GSS population based on age, gender, income, and belief that the government should reduce income differences. Results remain unchanged.

As Engelmann (2012) emphasizes, introducing a surplus term into the standard Fehr-Schmidt model makes the coefficients difficult to interpret. For example, suppose we control for the change in total surplus in equation (6). Then, the effects of $DiffIncome^1$, $DiffIncome^7$ and $DiffSurplus$ could be re-interpreted as the effects of $DiffIncome^1$, $DiffIncome^7$ and the total change in all *other* positions (as $\sum_p DiffIncome_p = DiffSurplus$). Put differently, holding $DiffSurplus$ constant (as we implicitly do when we control for it) while increasing $DiffIncome^p$ requires that some $DiffIncome^{p'}$ for $p' \neq p$ must decrease. Nonetheless, for the sake of completeness, we include the change in surplus in Appendix Table A.3. Our coefficients of interests remain unchanged. In particular, the coefficient on $DiffIncome^1$ barely falls, suggesting that little of the observed preference for raising the income of the poorest person is explained by a desire to raise total surplus.

In a similar vein, we include the difference in Gini coefficients between the two distributions as a control in Appendix Table A.4, to ensure that the emphasis we document over $DiffIncome^1$, $DiffIncome^7$, and $DiffIncome^{+1}$ are distinguishable from a general distaste for inequality. Again, we find that the coefficients on our variables of interest are largely unchanged.[14]

In Table 3 we repeat our analyses from Table 2, replacing the $DiffIncome$ variables with $SignIncome$ variables (recall, a dummy for whether a given value in Distribution B is larger

[13]One further question-order concern is subject fatigue over the course of the 10 iterations. We generate very similar point estimates when we restrict our analyses to the first or last two iterations for each subject which indicates that in practice choices are quite stable across the experiment.

[14]Specifications simultaneously controlling for Gini and surplus lead to near-identical results on the variables of interest.

than that in A).[15] The results are qualitatively similar, but are more readily interpretable. Consider the estimates in column (1). The coefficients on $SignIncome^1$, $SignIncome^7$, and $SignIncome^{+1}$ are 0.299, -0.107, and -0.0873 respectively (all significant at the one-percent level), whereas the coefficient on $SignIncome^{+2}$ is very close to zero. These results indicate that a subject is nearly thirty percentage points more likely to select Society B if the income of the poorest individual in that distribution is higher than the income of the poorest individual in Society A. The coefficient estimates also indicate a significant concern for reducing the incomes of individuals in the highest position and those in the position immediately above the subject's own. These latter two effects are of comparable magnitudes, and about a third as large as the effect of the poorest individual's income.

5.3 Implications for models of social preferences and redistribution

The implications of our results for models of social preferences can most easily be discerned in the context of the Fehr and Schmidt (1999). In the original formulation, the fixed above- and below-weights are inconsistent with the very different importance that our subjects place on the topmost and bottom-most positions in society. Rather, choices are better characterized by a simple and parsimonious parameterization of a more general Fehr-Schmidt style of model, presented in Section 2, in which an individuals' concern for payoffs diminish rapidly once one moves any distance from the best- and worst-off members of society. Additionally, our findings place a particular emphasis on immediate neighbors' incomes.

We compare the explanatory power of the Fehr-Schmidt model versus our specification in Table A.5. Column (1) repeats our main specification from Table 2, column (1), whereas column (2) shows a specification which includes instead the Fehr-Schmidt variables reflecting advantageous and disadvantageous inequality. Column (3) includes both sets of variables. A comparison of columns (1) and (3) shows that the variation explained (as captured by the R-squared) of our main specification is largely unchanged by the addition of the Fehr-Schmidt coefficients, and the point estimates on our variables of interest are similarly unaffected. By contrast, the fit of the Fehr-Schmidt model in column (2) is substantially improved by the introduction of our variables of interest, and their introduction sharply reduces the explanatory power of both Fehr-Schmidt variables.

Our findings also may be interpreted as reflecting features of other canonical social preference formulations. Most notably, as with Engelmann and Strobel (2004), we find that our subjects are both attentive to the payoffs of the worst-off individual, and also attend to efficiency, in the sense that surplus is predictive of subjects' choices.

[15] When we include the difference- and sign-based measures in the same specification, both groups of variables remain significant at least at the 5% level.

Our subjects' decisions are, broadly speaking, also consistent with models in which individuals put negative weight on individuals' incomes above their own, holding surplus constant (e.g., Charness and Rabin (2002) (with a negative weight on other's payoff when an individual is behind), Bolton and Ockenfels (2000), and Fehr and Schmidt (1999) all have this feature). However, our setting (in contrast to the experimental evidence presented in, for example, Charness and Rabin (2002)), allows us to discern that such concerns are very much concentrated at the high end of the distribution. This feature of our analysis may be particularly salient for (or indeed a direct result of) the runaway incomes at the top end of the distribution in the U.S. (see Piketty et al. (2018)).

Finally, our results may ultimately be a useful input into models of taxation and redistribution, which have tended to focus on the assumption that individuals maximize their own payoffs in voting for a proportional tax (e.g., Meltzer and Richard (1981) and its antecedent Romer (1975)). While our experiment suggests that individuals have a wider set of distributional concerns, by construction we do not allow for a clear weighting of payoffs to oneself versus distributional concerns, as we hold own-payoff constant. But our findings do suggest the basic intuition that, holding concerns for own-payoff constant, subjects' choices indicate a highly progressive tax policy, which may serve to help the worst-off, financed by taxes on the best-off members of society. Indeed, our findings indicate that tax on the wealthiest serves not simply as a source of revenue, but also is a direct input into the utility of less well-off members of society. The one paper that we know of which examines the implications of distributional preferences for voting on redistributive policies is Höchtl et al. (2012), which does so experimentally in a much-simplified framework. We see this as a very fruitful direction for further research, potentially guided by our results and research more broadly on distributional preferences.

Naturally, voting is a decision that is distinct from the social planner choices our subjects make in the experiments we describe above. We attempt to bridge this gap in a "democracy" treatment that we describe in Section 6, along with a number of other treatments we implemented to explore interpretations and robustness of our results.

5.4 Heterogeneity in Distributional Preferences

In our next set of analyses, we explore the extent to which our estimated effects from equation (6) vary systematically with political or self-stated distributional preferences, using our main sample. In columns (1) and (2) of Table 4, we compare the decisions of self-identified Democrats and Republicans (many subjects identified as independents, which is why the total sample is smaller). We conjecture that, given the Republican Party platform in recent decades of lowering taxes, its supporters will be less apt to choose distributions that reduce inequality. Consistent with this view, we find that Republicans are less likely to choose distributions with lower incomes in the top position (i.e., the coefficient on $DiffIncome^7$ is less negative in column (2) than in column (1)). Similarly, Republicans are less likely to select distributions with higher incomes in the lowest position. But we observe no difference between the two subsamples in their attitudes toward incomes of those directly above them — in both instances we observe a strong local competition effect.

In columns (3) and (4) we divide the sample based on responses to the question, "Do you feel that the distribution of income and wealth in the U.S. today is fair or should be more evenly distributed among a larger portion of the population?" Those in column (3) take the more redistributive position that income should be more evenly divided, whereas those in column (4) take the position that redistribution is not needed. In general, this cut of the data reveals starker differences in preferences than we saw in the first two columns. The coefficient on the poorest person's income is 67 percent larger for those in column (3) than in column (4). Even more striking differences between the two groups emerge in how they view the income of the richest person. For those who feel no more redistribution is needed in the U.S., the income of the richest person has no predictive power over which distribution is chosen (though the coefficient is negative). The corresponding coefficient for those who feel more redistribution is needed is over seven times larger in magnitude, negative and highly significant. The fact that the two groups differ far more on their views on the incomes of the rich than their views on the incomes of the poor may reflect the oft-stated conservative principle that inequality *per se* is not a concern relative to ensuring decent opportunities for the poor.[16]

Interestingly, however, despite these disparate views on incomes at the tails of the distribution, we see no substantial difference in the coefficients on $DiffIncome^{+1}$ and $DiffIncome^{+2}$. Overall, we take these findings as an indication of that the local competition effect may be quite distinct from preferences toward income inequality in general, which tend to focus more on the best- and worst-off members of society. While concern for the top- and bottom-levels of income appears correlated in the expected manner with self-identified ideology, concern for local disadvantageous inequality is present for both Democrats and Republicans and is present regardless of views about the fairness of the U.S. income distribution.[17]

[16]See, e.g., Mankiw (2013), who writes, "To the extent that our society deviates from the ideal of equality of opportunity, it is probably best to focus our attention on the left tail of the income distribution than on the right tail."

[17]In Appendix Table A.6 we show the same heterogeneity analysis for the session in which own income was varied. We find broadly similar patterns: those who feel that income should be more evenly distributed in the U.S. are also willing to give up more of their money to help the poorest person or to lower the income of the richest person, relative to those who feel the current U.S.

Finally, in Appendix Table A.7, we explore whether the patterns we report in our main result differ by subject income, age, gender, or education. For the first two subject characteristics, we split the sample at the median. Across all columns, we find remarkable stability in the local competition effect—the coefficients in each pair are near-identical. While we find some differences in top-most and bottom-most inequality aversion (e.g., women exhibit somewhat greater top-most inequality aversion than men), the overall patterns are quite consistent across sample splits.

6 Results from companion experiments: further robustness and extensions

So far we have shown that the results from our main experimental sessions are robust to a wide range of specifications. We now document the results from the companion experiments mentioned earlier to assess the robustness of our findings to changing various aspects of the experimental design.

6.1 Description of further experiments

After our two main experimental sessions, we conducted nine additional experiments, each of which significantly changed some property of the original experiment (listed in the chronological order in which they were implemented).

1. The OV ("own variation") experiment. This experiment allows the subject's own income to vary between distributions A and B. However, in both distributions he is in the same rank.

2. The NP ("nine person") experiment. This experiment tests whether the main results are robust to increasing the number of members in each distribution. For this experiment, we begin with an interval of $10,000-$20,000, with the increment increasing by $4,000 for each interval, so that $202,000-$244,000 is the highest interval.

3. The HI ('"high inequality") experiment. In the "high inequality" version, the lowest income range was $10,000-$15,000, with the income ranges increasing by $10,000 at each increment (so the top range was $190,000-$255,000).

distribution is fair. We view these results as reassuring checks on the validity of our main results. In general, we find similar patterns across all of our companion experiments, but in the interest of space do not report these results.

4. The VI ("very high inequality") experiment. In the "very high inequality" version, the income ranges increased by \$15,000 at each increment (so the top range was \$265,000-\$360,000).

5. The AF ("alternative framing") experiment. In this version, we provide an alternative presentation of the data, with the subject's own income presented to the far left of each panel in every decision. See Figure 2. The purpose of this version of the experiment was specifically to explore whether the local competition effect was attenuated by drawing subjects' attention *away* from the area of the graph immediately around their own incomes.

6. The RS ("real stakes") experiment. In this version, subjects were informed that, with 10 percent probability, one of their rounds would be implemented for a scaled down version (with each value divided by $10,000$) of the chosen income distribution.

7. The D ("democracy") experiment. The choice between societies in this variant was presented to respondents as the outcome of a democratic process, in which they were to vote on their desired outcome. It uses the percent-differences distributions.

8. The LM ("labeled mean") experiment. This version was identical to our main experiment except that the mean of the distribution was labeled on each graph. The purpose of this experiment is to explore the consequences of making another income level – the mean – particularly salient.

9. The H ("histogram") experiment. This variant of our experiment differs most significantly from our main design. Instead of presenting subjects with a distribution of incomes, the histogram treatment presents a histogram capturing the fraction of individuals in income bins of \$10,000, topcoded at \$250,000. This version is meant to explore the extent to which some of the patterns we observe in our 7-9 person experiments carry over to societies that are more reflective of those that our subjects inhabit.

In total, we have conducted eleven experiments for this project, and all but the histogram treatment can be directly compared as their coefficients have the same units. Because the histogram treatment departs so substantially from our the presentation in our main experiment, we provide a separate discussion of these results in Section 6.3.

6.2 Results from companion experiments

In Figure 5 we show the estimated parameters (along with confidence intervals) for each variant described above (excepting the histogram treatment), to capture bottom- and top-most inequality aversion as well as locally competitive preferences, based on the specification from column (1) of Table 2. For completeness, a full set of results paralleling those presented in Table 2 are available for each additional session in a series of tables in Appendix C. (These include the two sessions from our main sample, but separated by session, to ensure that the results we report in Table 2 are not driven by just one of the two sessions.)

Across all sessions, the coefficients on $DiffIncome^1$ and $DiffIncome^7$ (or, in the case of the NP, $DiffIncome^9$, which captures the highest incomes in that treatment) are of consistent sign. These results indicate that aversion to inequality at both the high and low extremes of the income distribution is largely robust to the type of distribution, its presentation, as well as the introduction of payoff consequences for subjects' choices.

Our estimate of local competitiveness, as captured by $\beta_2 - \beta_1$, is positive across almost all variants, and is significantly different from zero in six of the ten sessions as well (it is statistically significant at the ninety-percent level in seven of the ten). Of particular note is the robustness of the effect in the "alternative framing" variant, as that session made it more difficult for the subject to compare her own income to her neighbors.

The local-competition effect is indistinguishable from zero in the real-stakes treatment (for which both coefficients are quite close to zero) and the variant in which the mean is labeled.(For each of these variants, we can reject that the difference in coefficients is equal to that of our main experiments at the 1% level.).

It may be instructive to consider in greater detail the circumstances in which the "local competition" is not readily discernible. We begin with the real-stakes treatment. While in this variant individuals are still more likely to choose the distribution with the lower income for the person directly above, this result is no longer significant, and when we compare this coefficient to that of the person two positions above, the difference, while still of the predicted sign, is smaller in magnitude relative to our main experiment, and is statistically insignificant.

The prior literature provides little guidance on the difference between the results in our real-stakes and hypothetical treatments. Camerer and Hogarth (1999), in particular, provide a meta-analysis of 74 studies that have no, low or high-powered incentives. They find that the effect of real stakes depends on the experimental task.[18] None of these experiments,

[18]Some more recent experiments since the meta-analysis was published focus specifically on comparing behavior with and without payoff consequences, but again find mixed results. Further, none of these recent studies invokes the sort of distributive concerns that are our focus in this

however, concern the types of distributive principles that we explore here. Moreover, we are also intentionally evoking the *actual* income distribution, which most experiments do not do. The only redistribution experiment we know of that compares hypothetical and real stakes is Charité et al. (2015), who find similar results in a modified dictator game with and without real stakes, though that experiment did not try to frame outcomes in terms of actual, real-world income distributions.

The differing results across real stakes versus hypothetical treatments potentially raise deeper methodological questions on the measurement of distributional preferences in lab experiments. In particular, we are interested in studying distributional preferences over total income or wealth, so that it is naturally impossible to implement subjects' choices in practice. As a result, when we impose payoff consequences we may substantively shift the distributive principles that subjects invoke in making their decisions. That is, subjects may have in mind the fairness principles toward a *society's* income distribution overall when making choices without direct payoff consequences. But when they are told that some specific, rather arbitrary handful of actual people (those MTurk workers we are rewarding with $\frac{1}{10{,}000}$ of these "real-world" income values) will experience these payoffs, they may invoke different principles.[19]

Next, we examine the labeled-mean treatment. One natural consideration in explaining the disappearance of the local competition effect is that average income becomes a particularly salient point of comparison when highlighted, and thus may draw attention away from other relative ranking benchmarks. Consistent with this interpretation, the sign on income surplus (which is simply a constant times the mean) becomes *negative* when mean income is labeled, suggesting that, perhaps due to experimenter demand effects and perhaps due to salience, respondents use it as a benchmark for their own incomes when it is highlighted (see Appendix Table C.10).

---

paper. See, for example, Etchart-Vincent and l'Haridon (2011) and Beattie and Loomes (1997).

[19]One particularly salient feature of this difference is that marginal utility of income enters more into the real-world framing. In the classic inequality-aversion set up, utility is linear in own-income, an approximation that is likely innocuous for the small-stakes settings in which it is typically tested. If subjects were highly sensitive to concerns about the diminishing marginal utility of a dollar, we might expect that the coefficients on the tail incomes in the "real stakes" version to be smaller in magnitude than in the other sessions (since the amount of money involved would have trivial effects on the marginal utility of income in the "real stakes" version, but potentially large effects in the others). Comparing the "real stakes" session to the other versions in Figure 5 the coefficients on the top and bottom incomes are generally quite similar (the coefficients found in the "real stakes" version are roughly at the midpoint of the range formed by the full set of sessions). So, taken literally, our results seem to suggest that these concerns were not paramount to our subjects. However, we find this distinction (between true inequality aversion and beliefs about the diminishing marginal utility of money) to be a very interesting question for future work.

This "average as benchmark" interpretation is further bolstered by an analysis of responses to the open-ended questions for this treatment. These open-ended questions were included in three of the treatments described above: the labeled mean, real stakes, and democracy treatments. In Appendix Table A.8 we present the most common two-word (bigram) phrases that appear in these open-ended responses (three-word phrases, or trigrams, turned out not to be any more informative), disaggregated by treatment.[20]

For both the real stakes and democracy treatments, the most common bigrams all relate to higher/highest and lower/lowest income. These appear to reflect concerns for the best- and worst-off members of society, and also potentially concern for individuals just-higher in the distribution.[21] By contrast, when we label the mean, the most common (by a very wide margin) is "average incom." Furthermore, when one looks at the full comments that mention average income in this treatment, they tend to describe a desire to be as far above average (or if below, as close to the average) as possible. For example, to mention just a few specific comments: "I compared average income with mine to make my choice"; "I wanted to have the highest income relative to others. I tried to make choices where I was richest relative to the average income"; "I either chose to be as much above average as possible, or if lower than average, I chose the selection closest to the average."

In summary, the labeled-mean treatment obscures the concern for lowering the income of the person just above, but has limited effect on concern for those at the bottom and the top. The results from the labeled-mean variant along with all other variants suggest that, beyond focusing on the top and bottom, respondents appear to naturally consider the incomes of people closer to them. When the mean is not emphasized (as in most of our variants), they appear to focus on those just above them. When the mean is emphasized, that value becomes the "local" benchmark.

[20]We use the "tm" package in R to process the text of the responses to this question. We convert all text to lowercase, strip punctuation and common English stopwords, and stem words with a Porter stemmer. We then take all 2-word (bigram) and 3-word (trigram) sequences in the remaining text, and calculate frequencies across subject responses.

[21]Bigrams are less suited to capturing subjects' expressions of local competition, which is a more complicated notion of distributional preferences. But one indication of such concerns may be found among subjects who used the relatively common bigram "higher income" (38 mentions in the Real Stakes version, 14 in the Democracy one). Some explicitly noted a concern for the incomes of those just above them (e.g., "I also Looked at the difference between each income level esp between the level that was given to me and the next tier up from that."). Much more commonly, however, the implication was ambiguous on the weightings given to higher incomes (e.g., I chose my distributions by looking at...how far I was from the higher income people on the graph.")

6.3 Using histograms to represent income distributions

All of our distributions up to this point have been illustrated via bar graphs. This choice has the advantage of being the easiest for subjects to understand. Students are often introduced to bar graphs as early as kindergarten (Cooper and Shore, 2010), and even some first-graders naturally create bar graphs on their own when asked to describe a particular characteristic of an assortment of items (English, 2013).

The final variant of our experiment employs histograms instead, which we do for two reasons. First, histograms are potentially the most natural way to illustrate more continuous distributions such as, say, the income distribution of the over 100 million households in the US. As such, it helps us move from distributions of seven or nine people ("large" relative to the past literature, but small in absolute number) to one meant to proxy the full population. Second, the histogram illustration is profoundly different from a bar-graph representation, so it helps test robustness to presentation.

Appendix Figure B.1 shows the typical choice that confronts a subject in this variant of the experiment, and Appendix B provides the instructions given to subjects in this variant. Bin choice is a key parameter for a histogram, and we chose the width and number of bins so as not to have a large mass in the final bin. Given the skew of the U.S. household income distribution, some mass in the final bin is unavoidable.

Of course, the bar heights in a histogram represent *frequencies* and the x-axis represents income values (whereas in a bar graph, the income values are on the y-axis). Thus, instead of asking whether subjects prefer to lower the top income or raise the bottom income, we are asking whether they prefer a society with *fewer* rich or fewer poor people. Top- and bottom-most inequality aversion would predict a negative coefficient on the height of the top-most income category, as well as a *negative* coefficient for the bottom-most group (in contrast with the previous variants). As the coefficients for this variant are in different units than in all the others (and in fact the predicted sign differs for one), we do not include the coefficients for this variant in Figure 5.

The results in Table A.10 are consistent with our predictions: the coefficient on the differences (again, in *frequency* in this case, not in income as in all other variants of the experiment) are negative for both the top and bottom income group. Our subjects prefer to shrink the mass at the tails on both sides of the histogram, consistent with top- and bottom-most inequality aversion. Similarly, people prefer to have fewer people in the income group right above them (relative to the size of the group two levels above them).

While in all cases the sign of the coefficients are in the expected direction for someone who is averse to top- and bottom-most inequality and local competition, the coefficient for the bottom-most result is not significant.

At first, the lack of significance for bottom-most inequality aversion might be surprising, given the robustness of this result in all other variants. However, the histogram version of our experiment may be more subject to misinterpretation by subjects. Indeed, there is a large pedagogical literature documenting the difficulty most people have comprehending histograms. The common problem they exhibit is mistaking the height of the bars for values (in our case, income) instead of frequencies. As Whitaker and Jacobbe (2017) write in a review of the evidence: "While bar graphs are generally understood by and can be used by students, studies of student understanding of histograms have revealed numerous misunderstandings and areas of confusion....Among the most widely reported misunderstanding about bar graphs and histograms is that the variability in the data is represented in a histogram by variability in the *heights* of the bars [emphasis added]." There is a vast literature in math and statistics education documenting this error: it exists even for college students (in fact, even for grades 4–12 math *teachers*), and persist on post-tests even after instruction.[22] As noted earlier, this problem does not exist for bar graphs, as even young children are generally able to interpret them correctly. Scholars of pedagogy speculate that students' inherent comfort with bar graphs is in part responsible for their confusion with histograms, and they thus interpret the latter as the former. In Appendix Table A.9 we show that our sample very likely suffers from the same confusion: in column (2), the coefficient on the height of their *own* bar in the histogram is very highly significant, suggesting they incorrectly interpret bar height as income (as there is no obvious reason individuals would want their income group to be *large*, as would be the implication if they interpret the figure as a histogram).

For someone who is top-most inequality averse, interpreting the graph correctly as a histogram or incorrectly as a bar graph both produce the same response: she would choose the distribution with a shorter bar in the highest position (either wanting fewer people in the highest income group or a highest income group with a lower average income). But for someone who is bottom-most inequality averse, the choice would depend on whether she interprets the graph as a histogram or bar graph: if a histogram, she would want fewer people in the worst-off category (choose a lower bar height), but if a bar graph she would want to raise the income of the worst-off group (choose a higher bar height). In Columns (3) and (4), we limit the sample to those whose choices imply a coefficient on own bar-height close to zero, which we use as a (very rough) proxy for understanding the histogram.[23]

[22] See delMas et al. (2005), Shore and Cooper (2009) and citations therein.

[23] We run the following regression separately for each individual: $Choose^B_{ik} = \beta_1 SignFrequency^1_{ik} + \beta_2 SignFrequency^{26}_{ik} + \beta_3 SignFrequency^{OWN}_{ik} + e_{ik}$. Appendix Figure B.2 displays the distribution of estimated β_3's. For this sample, the coefficient on the top frequency remains negative, and now the negative coefficients implied by local competition and bottom-most inequality aversion are also significant.

The coefficient on the frequency for the lowest-income group remains negative, grows in magnitude and becomes marginally significant. The effects for top-most inequality and local competition retain their expected signs and remain significant.

7 Conclusion

In this paper, we study distributional preferences in "large" (seven to nine person) groups, using a series of online experiments conducted via Mechanical Turk. We find a very robust and consistent emphasis on reducing extreme inequality: consistent with Rawlsian preferences, subjects are much more likely to select an income distribution that leads to a higher income for the poorest individual. More novel, we find a robust preference for distributions that, all else equal, have lower incomes for the richest individual, and also (in most experimental sessions) a preference for reducing the incomes of individuals directly above the subject.

Our experiments were quite abstract in the sense that there was no mention of how Society A might be transformed into Society B. Future work may wish to emphasize the role that taxes and transfers would necessarily play in comparing two hypothetical societies. Recent work has found that subjects often react differently to variation they were told was exogenous versus variation they were told was driven by taxation. For example, Kessler and Norton (2016) find that labor supply (in the form of real effort) drops more when subjects are told that a tax has been taken out of an experimental wage than when the wage is merely lowered (even though the change in the effective wage was the same).

The abstraction of our experiment also begs the question of how the notion of "bottom-most," "top-most," and "local" inequality aversion translate into the incomes distributions that individuals may encounter in practice. We offer several speculative thoughts on the matter. First, one way of conceiving of the real-world analog to our seven person society may be to think of each individual as representing groups in the income distribution. The top of the distribution may correspond most readily to the "one percent" and calls for taxing this group as reflecting our subjects' desires to reduce incomes at the high end. One may also think of our seven person society as reflecting the incomes of those an individual can readily observe and most readily influence her beliefs on redistribution. Domenech (2020), for example, examines the extent to which neighborhood inequality affects survey respondents' beliefs about the overall income distribution in Spain; as in our experiment, this study emphasizes that certain incomes and comparators may be more salient than others.

While our work is motivated by a desire to better understand attitudes toward income inequality for the U.S. overall, the decisions confronting our subjects may also be relevant for inequalities in more intimate groups. For example, our results may be applied to understanding (and devising empirical tests for) attitudes toward pay inequalities within companies or other organizations.

We hope that our experimental findings can provide some guidance on how individuals weight income gaps between themselves and others in particular positions in a given distribution. While our results, taken from a combination of real stakes and hypothetical experiments, should be interpreted with caution, we hope that it will spur further work to enrich our understanding of how individuals conceive of inequality in larger groups that have more direct relevance for the types of redistributive decisions confronted by society.

References

Ackert, Lucy F, Jorge Martinez-Vasquez, and Mark Rider, "Social preferences and tax policy design: some experimental evidence," *Economic Inquiry*, 2007, *45* (3), 487–501.

Almås, Ingvild, Alexander W Cappelen, Erik Ø Sørensen, and Bertil Tungodden, "Fairness and the development of inequality acceptance," *Science*, 2010, *328* (5982), 1176–1178.

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.

Beattie, Jane and Graham Loomes, "The impact of incentives upon risky choice experiments," *Journal of Risk and Uncertainty*, 1997, *14* (2), 155–168.

Beckman, Steven R, John P Formby, and W James Smith, "Efficiency, equity and democracy: Experimental evidence on Okuns leaky bucket," in "Inequality, welfare and income distribution: Experimental approaches," Emerald Group Publishing Limited, 2004, pp. 17–42.

Bernardi, Richard A, "Associations between Hofstedes cultural constructs and social desirability response bias," *Journal of Business Ethics*, 2006, *65* (1), 43–53.

Bolton, Gary E and Axel Ockenfels, "ERC: A Theory of Equity, Reciprocity, and Competition," *American Economic Review*, 2000, pp. 166–193.

Camerer, Colin F and Robin M Hogarth, "The Effects of Financial Incentives in Experiments: A Review and Capital-Labor-Production Framework," *Journal of Risk and Uncertainty*, 1999, *19* (1), 7–42.

Cappelen, Alexander W, Trond Halvorsen, Erik Ø Sørensen, and Bertil Tungodden, "Face-saving or fair-minded: What motivates moral behavior?," *Journal of the European Economic Association*, 2017, *15* (3), 540–557.

Charité, Jimmy, Raymond Fisman, and Ilyana Kuziemko, "Reference Points and Redistributive Preferences: Experimental Evidence," Technical Report, National Bureau of Economic Research 2015.

Charness, Gary and Matthew Rabin, "Understanding Social Preferences with Simple Tests," *Quarterly Journal of Economics*, 2002, *117* (3), 817–869.

Cooper, Linda L. and Felice S. Shore, "The Effects of Data and Graph Type on Concepts and Visualizations of Variability," *Journal of Statistics Education*, 2010, *18* (2), null.

Dalton, Derek and Marc Ortegren, "Gender differences in ethics research: The importance of controlling for the social desirability response bias," *Journal of Business Ethics*, 2011, *103* (1), 73–93.

delMas, Robert, Joan Garfield, and A Ooms, "Using assessment items to study students difficulty reading and interpreting graphical representations of distributions," in "Fourth Forum on Statistical Reasoning, Thinking, and Literacy (SRTL-4)" 2005.

Domenech, Gerard, "Perceived Inequality and Preferences for Redistribution," 2020.

Durante, Ruben, Louis Putterman, and Joël Van der Weele, "Preferences for redistribution and perception of fairness: An experimental study," *Journal of the European Economic Association*, 2014, *12* (4), 1059–1086.

Engelmann, Dirk, "How not to extend models of inequality aversion," *Journal of Economic Behavior & Organization*, 2012, *81* (2), 599–605.

_ **and Martin Strobel**, "Inequality Aversion, Efficiency, and Maximin Preferences in Simple Distribution Experiments: Comment," *American Economic Review*, 2004, *96* (5), 1906–1911.

English, Lyn D., "Surviving an Avalanche of Data," *Teaching Children Mathematics*, 2013, *19* (6), 364–372.

Etchart-Vincent, Nathalie and Olivier l'Haridon, "Monetary incentives in the loss domain and behavior toward risk: An experimental comparison of three reward schemes including real losses," *Journal of risk and uncertainty*, 2011, *42* (1), 61–83.

Farhi, Emmanuel and Xavier Gabaix, "Optimal taxation with behavioral agents," Technical Report, National Bureau of Economic Research 2015.

Fehr, Ernst and Klaus M. Schmidt, "A Theory Of Fairness, Competition, and Cooperation," *Quarterly Journal of Economics*, 1999, *114* (3), 817–868.

Fisman, Raymond, 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.

Höchtl, Wolfgang, Rupert Sausgruber, and Jean-Robert Tyran, "Inequality aversion and voting on redistribution," *European economic review*, 2012, *56* (7), 1406–1421.

Holt, Charles A and Susan K Laury, "Risk aversion and incentive effects," *American Economic Review*, 2002, *92* (5), 1644–1655.

Kahneman, Daniel, Jack Knetsch, and Richard Thaler, "Fairness and the Assumptions of Economics," *Journal of Business*, 1986, *59* (4), S285–S300.

Kessler, Judd B and Michael I Norton, "Tax aversion in labor supply," *Journal of Economic Behavior & Organization*, 2016, *124*, 15–28.

Krawczyk, Michał, "A glimpse through the veil of ignorance: Equality of opportunity and support for redistribution," *Journal of Public Economics*, 2010, *94* (1), 131–141.

Kreuter, Frauke, Stanley Presser, and Roger Tourangeau, "Social Desirability Bias in CATI, IVR, and Web Surveys: The Effects of Mode and Question Sensitivity," *Public Opinion Quarterly*, 2008, *72* (5), 847–865.

Kuziemko, Ilyana, Michael I Norton, Emmanuel Saez, and Stefanie Stantcheva, "How elastic are preferences for redistribution? Evidence from randomized survey experiments," *The American Economic Review*, 2015, *105* (4), 1478–1508.

Lockwood, Benjamin B, "Optimal Income Taxation with Present Bias," 2016.

Luttmer, E.F.P., "Neighbors as negatives: Relative earnings and well-being," *The Quarterly Journal of Economics*, 2005, *120* (3), 963–1002.

Mankiw, N. Gregory, "Defending the One Percent," *Journal of Economic Perspectives*, September 2013, *27* (3), 21–34.

Meltzer, Allan H and Scott F Richard, "A Rational Theory of the Size of Government," *The Journal of Political Economy*, 1981, pp. 914–927.

Norton, Michael I and Dan Ariely, "Building a Better America—One Wealth Quintile at a Time," *Perspectives on Psychological Science*, 2011, *6* (1), 9–12.

Piketty, Thomas and Emmanuel Saez, "Income inequality in the United States, 1913–1998," *The Quarterly journal of economics*, 2003, *118* (1), 1–41.

_ , _ , and Gabriel Zucman, "Distributional national accounts: methods and estimates for the United States," *The Quarterly Journal of Economics*, 2018, *133* (2), 553–609.

Romer, Thomas, "Individual welfare, majority voting, and the properties of a linear income tax," *Journal of Public Economics*, 1975, *4* (2), 163–185.

Saez, Emmanuel and Stefanie Stantcheva, "Generalized social marginal welfare weights for optimal tax theory," *The American Economic Review*, 2016, *106* (1), 24–45.

Schumacher, Heiner, Iris Kesternich, Michael Kosfeld, and Joachim Winter, "One, Two, Many–Insensitivity to Group Size in Games with Concentrated Benefits and Dispersed Costs," *The Review of Economic Studies*, forthcoming.

Shore, F. and L. Cooper, "How do in-service teachers (grades 412) perceive variability in graphs?," 2009. Poster session presented at The Joint Statistical Meetings, Washington, DC.

Whitaker, Douglas and Tim Jacobbe, "Students' Understanding of Bar Graphs and Histograms: Results From the LOCUS Assessments," *Journal of Statistics Education*, 2017, *25* (2), 90–102.

Table 1: Basic summary statistics in MTurk sample compared to GSS sample

	(1) MTurk sample	(2) GSS sample
Female	0.436	0.545
	(0.496)	(0.498)
Age	32.42	47.46
	(11.00)	(17.24)
Has at least college education	0.572	0.316
	(0.495)	(0.465)
Household income	61.86	81.00
	(168.8)	(80.64)
Voted in last US presidential election	0.715	0.639
	(0.451)	(0.480)
Supports gov't redistribution (scale 1-7)	4.571	4.244
	(1.908)	(2.062)
Thinks hard work most important to get ahead	0.423	0.707
	(0.494)	(0.455)
Observations	6725	2538

Notes: Column 1 includes all sessions of the experiment. Only subjects who completed all 10 iterations are included. For re-takers, they are included only the first time they took the survey. Column 2 includes all adults in the 2014 General Social Survey (weighted with the provided individual-level weights). Income refers to household income (in units of $1,000).

Table 2: Bottom, Topmost and Local Inequality Aversion in the Main Sample

	Dep. var: Chose Distribution B over A			
	(1)	(2)	(3)	(4)
$DiffIncome^{+1}$ in B vs. A	-0.0371***	-0.0372***	-0.0372***	-0.0388***
	[0.00415]	[0.00416]	[0.00416]	[0.00438]
$DiffIncome^{+2}$ in B vs. A	-0.000224	-0.000329	-0.000328	-0.000106
	[0.00365]	[0.00365]	[0.00365]	[0.00382]
$DiffIncome^{BOTTOM}$ in B vs. A	0.195***	0.196***	0.196***	0.200***
	[0.00662]	[0.00662]	[0.00662]	[0.00690]
$DiffIncome^{TOP}$ in B vs. A	-0.0426***	-0.0425***	-0.0425***	-0.0442***
	[0.00407]	[0.00407]	[0.00407]	[0.00431]
Above Two/One Diff.	.03691***	.03691***	.03684***	.03866***
Question-Order FE	No	Yes	Yes	Yes
Position	No	No	Yes	Yes
Ex. short duration	No	No	No	Yes
Observations	14515	14515	14515	13229
R2	0.103	0.103	0.104	0.109

Notes: All regressions include robust standard errors clustered at the subject level. The sample in this table includes participants in the "Absolute Differences" and "Percentage Differences" experiments run in September 2013 (see text for details). Only subjects who completed all 10 iterations are included $^{*}p < 0.1,^{**} p < 0.05,^{***} p < 0.01$. In all specifications, monetary values are expressed in units of $\$10,000$ to make the table more readable. $DiffIncome^{+1}$ is the difference in income between Societies B and A for the individual in the position directly above the subject's own. $DiffIncome^{+2}$ is similarly defined for the individual two positions above the subject. $DiffIncome^{7}$ is the difference in income between Societies B and A for the richest (i.e., position 7) individual. $DiffIncome^{1}$ is similarly defined for the poorest individual.

Table 3: Bottom, Topmost and Local Inequality Aversion in the Main Sample, Sign-based Results

	Dep. var: Chose Distribution B over A			
	(1)	(2)	(3)	(4)
$SignIncome^{+1}$ in B vs. A	-0.0873***	-0.0876***	-0.0874***	-0.0879***
	[0.00911]	[0.00911]	[0.00911]	[0.00952]
$SignIncome^{+2}$ in B vs. A	-0.00199	-0.00185	-0.00177	0.000842
	[0.00864]	[0.00864]	[0.00864]	[0.00896]
$SignIncome^{BOTTOM}$ in B vs. A	0.299***	0.300***	0.300***	0.305***
	[0.0108]	[0.0108]	[0.0108]	[0.0111]
$SignIncome^{TOP}$ in B vs. A	-0.107***	-0.108***	-0.108***	-0.107***
	[0.0108]	[0.0108]	[0.0108]	[0.0112]
Above Two/One Diff.	.08534***	.08571***	.0856***	.0887***
Question-Order FE	No	Yes	Yes	Yes
Position	No	No	Yes	Yes
Ex. short duration	No	No	No	Yes
Observations	13605	13605	13605	12590
R2	0.109	0.110	0.110	0.113

Notes: All regressions include robust standard errors clustered at the subject level. The sample in this table includes participants in the "Absolute Differences" and "Percentage Differences" experiments run in September 2013 (see text for details). Only subjects who completed all 10 iterations are included $^{*}p < 0.1,^{**} p < 0.05,^{***} p < 0.01$. $SignIncome^{+1}$ is an indicator variable denoting that the income of the person in the position directly above the subject is higher in Society B than in Society A. $SignIncome^{+2}$ is similarly defined for the individual two positions above the subject. $SignIncome^{1}$ is an indicator variable denoting that the poorest individual (i.e., position 1) has a higher income in Society B than in Society A. $SignIncome^{7}$ is similarly defined for the richest (i.e., position 7) individual.

Table 4: Heterogeneity in Inequality Aversion

	Dep. var: Chose Second Distribution			
	(1) Democrat	(2) Republican	(3) Not Fair Distribution	(4) Fair Distribution
$DiffIncome^{+1}$ in B vs. A	-0.0310***	-0.0503***	-0.0374***	-0.0330***
	[0.00631]	[0.00927]	[0.00434]	[0.00726]
$DiffIncome^{+2}$ in B vs. A	0.00232	0.00860	-0.00407	0.0108
	[0.00598]	[0.00855]	[0.00411]	[0.00683]
$DiffIncome^{BOTTOM}$ in B vs. A	0.195***	0.145***	0.217***	0.130***
	[0.00887]	[0.0142]	[0.00608]	[0.0108]
$DiffIncome^{TOP}$ in B vs. A	-0.0517***	-0.0305***	-0.0569***	-0.00785
	[0.00530]	[0.00755]	[0.00363]	[0.00590]
Above Two/One Diff.	.03328***	.05893***	.03334***	.04383***
Question-Order FE	No	No	No	No
Position	No	No	No	No
Ex. short duration	No	No	No	No
Observations	5183	2293	10521	3866
Coef. $DiffIncome^{BOTTOM}$ Diff		.04973**		.08773***
Coef. $DiffIncome^{TOP}$ Diff		-.02124*		-.04902***
Coef. Above Two/One Diff. Diff		.02565		.01049

Notes: All regressions use robust standard errors clustered at the subject level. The sample in this table includes participants in the "Absolute Differences" and "Percentage Differences" experiments run in September 2013 (see text for details). Only subjects who completed all 10 iterations are included $^{*}p < 0.1, ^{**}p < 0.05, ^{***}p < 0.01$. In all specifications, monetary values are expressed in units of $\$10,000$ to make the table more readable. In columns (1) and (2) regressions are run separately on the subsamples of, respectively, self-identified Democrats and Republicans. In columns (3) and (4) we divide the sample based on responses to the question, "Do you feel that the distribution of income and wealth in the US today is fair or should be more evenly distributed among a larger portion of the population?" $DiffIncome^{+1}$ is the difference in income between Societies B and A for the individual in the position directly above the subject's own. $DiffIncome^{+2}$ is similarly defined for the individual two positions above the subject. $DiffIncome^{1}$ is the difference in income between Societies B and A for the poorest (i.e., position 1) individual. $DiffIncome^{7}$ is similarly defined for the richest individual.

Figure 1: Standard representation of the choice question in the experiment

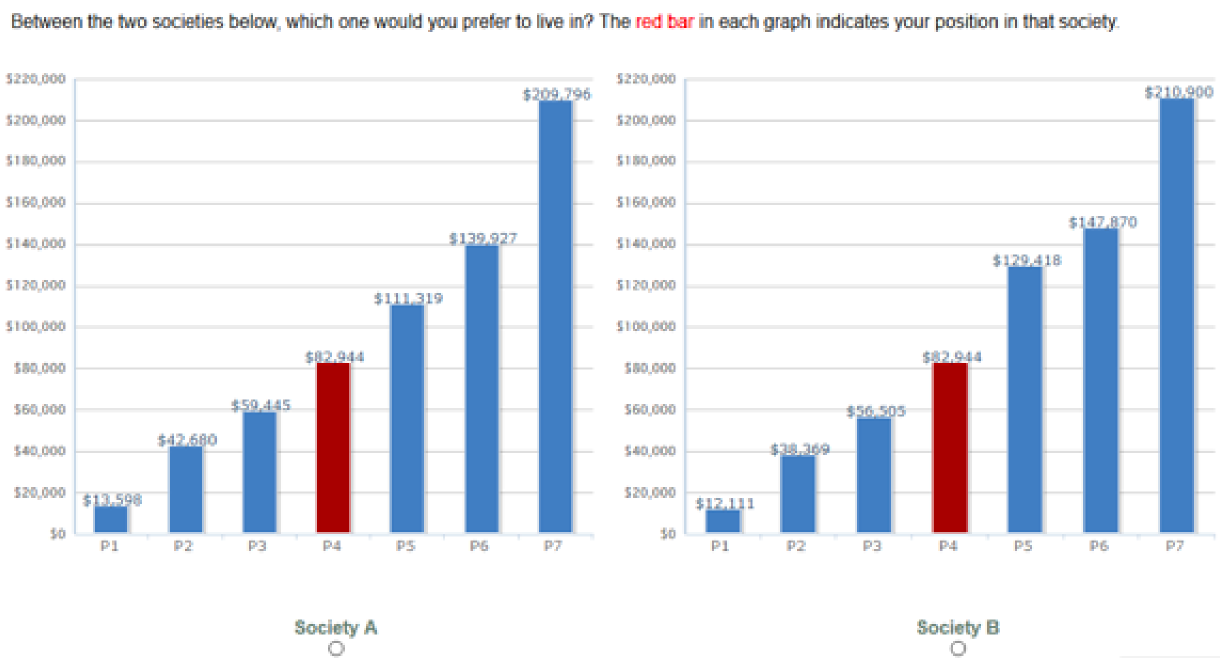

Figure 2: Alternative graphical representation of choice question

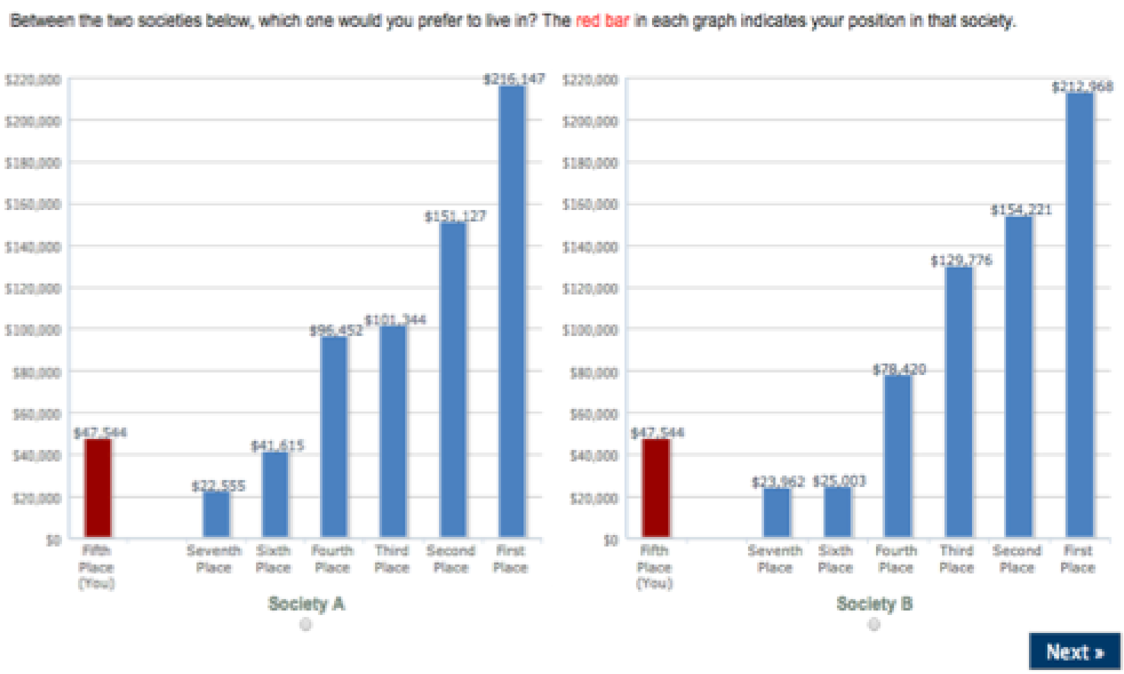

Figure 3: Effect of higher income in position q on subject's propensity to choose distribution B over A

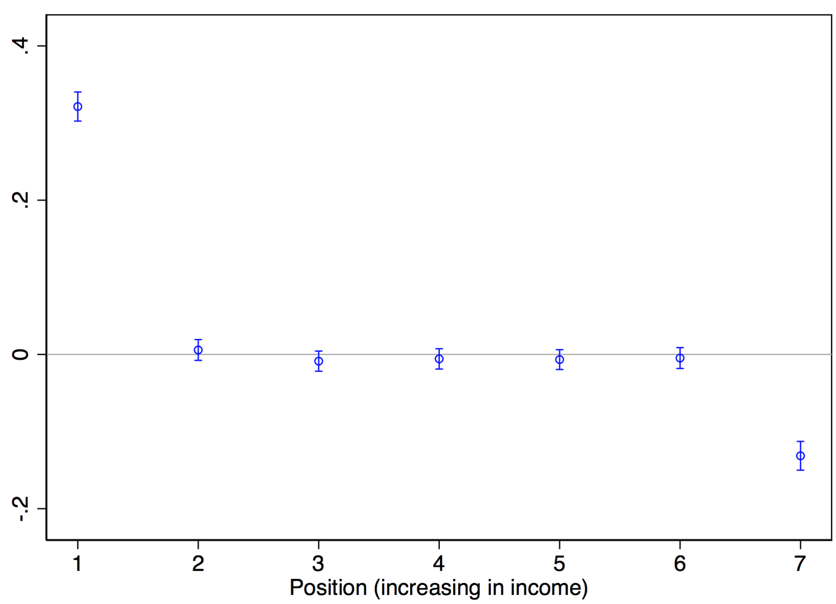

Note: The coefficients plotted for each graph are generated by regressing $Choose^B_{ik} = \alpha + \sum_{q=1}^{7} \lambda_q SignIncome_{q,ik} + P_{ik} + \epsilon_{ik}$, where i indexes the subject and k a particular iteration of the experiment (i.e., equation (4) in the text). The 95-percent confidence intervals plotted are based on standard errors clustered by subject.

Figure 4: Effect of higher income in position q on subject's propensity to choose distribution B over A, separately by position p of subject herself

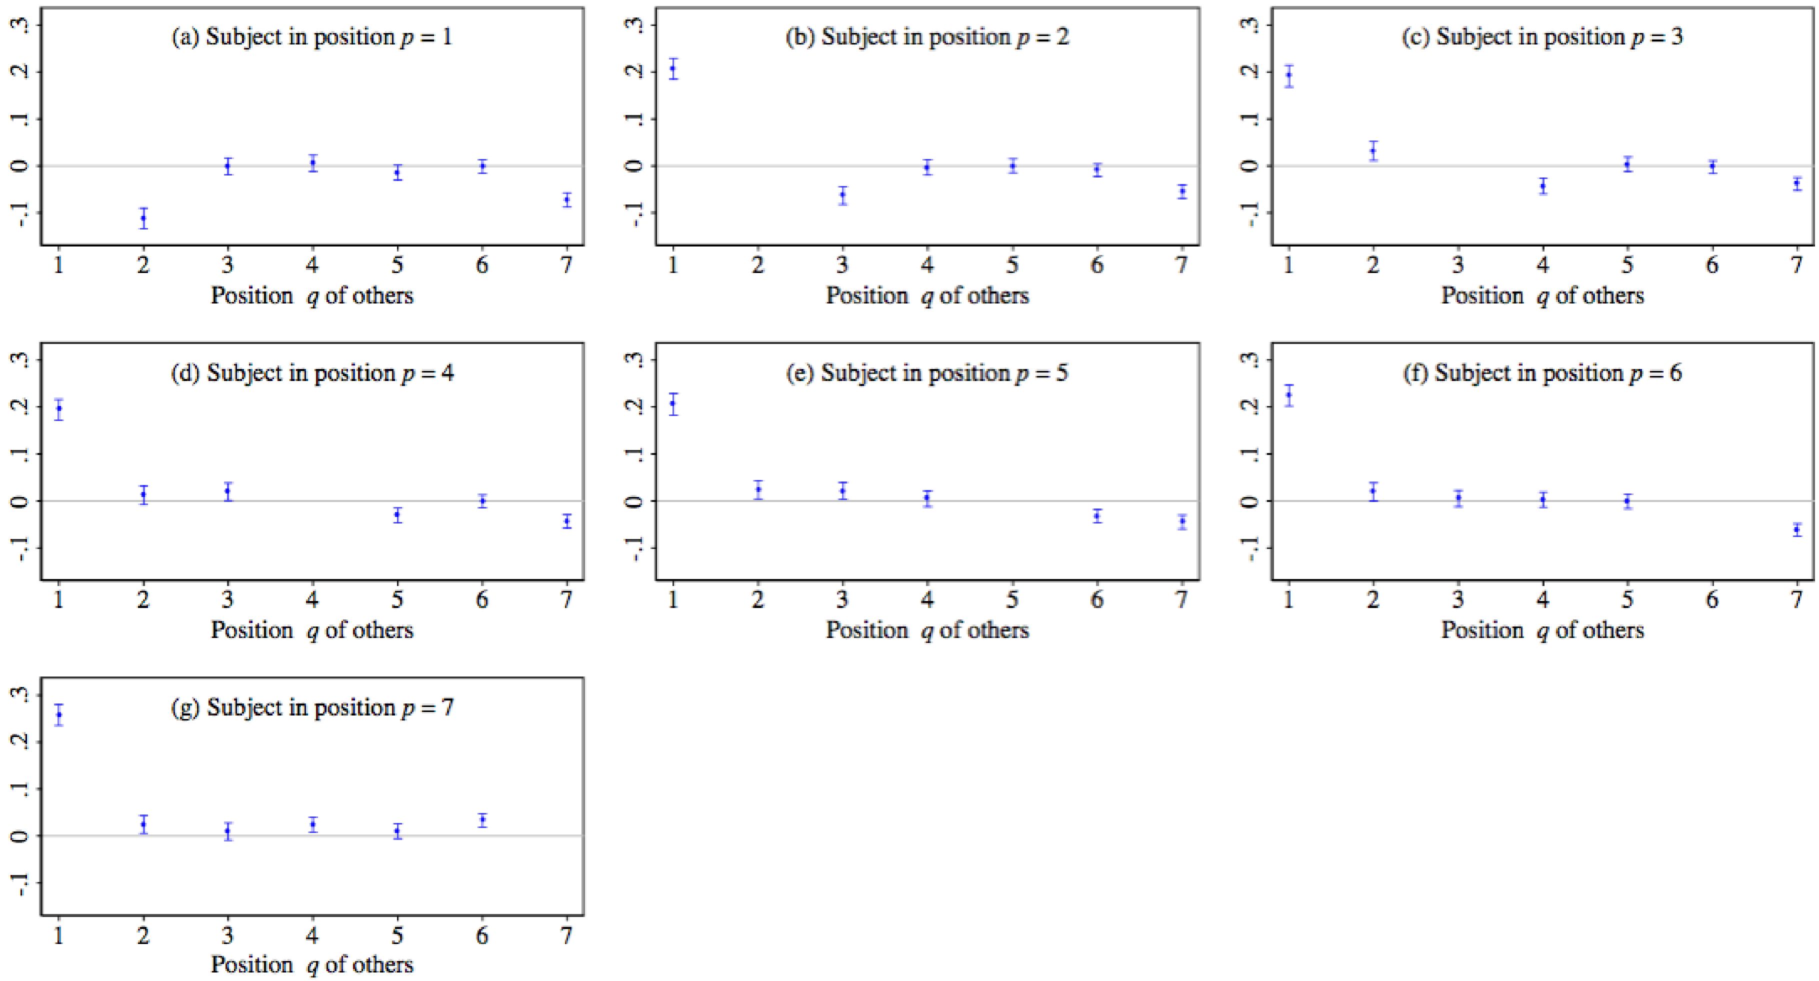

Note: The coefficients plotted for each graph are generated by regressing (for each position p) $Choice^B_{ik} = \alpha + \sum_{q \neq p} \eta^p_q SignIncome_{q,ik} + \epsilon_{ik}$ where i indexes the subject and k a particular iteration of the experiment (i.e., equation (5) in the text). The 95-percent confidence intervals plotted are based on standard errors clustered by subject.

Figure 5: Results Across Experimental Treatments

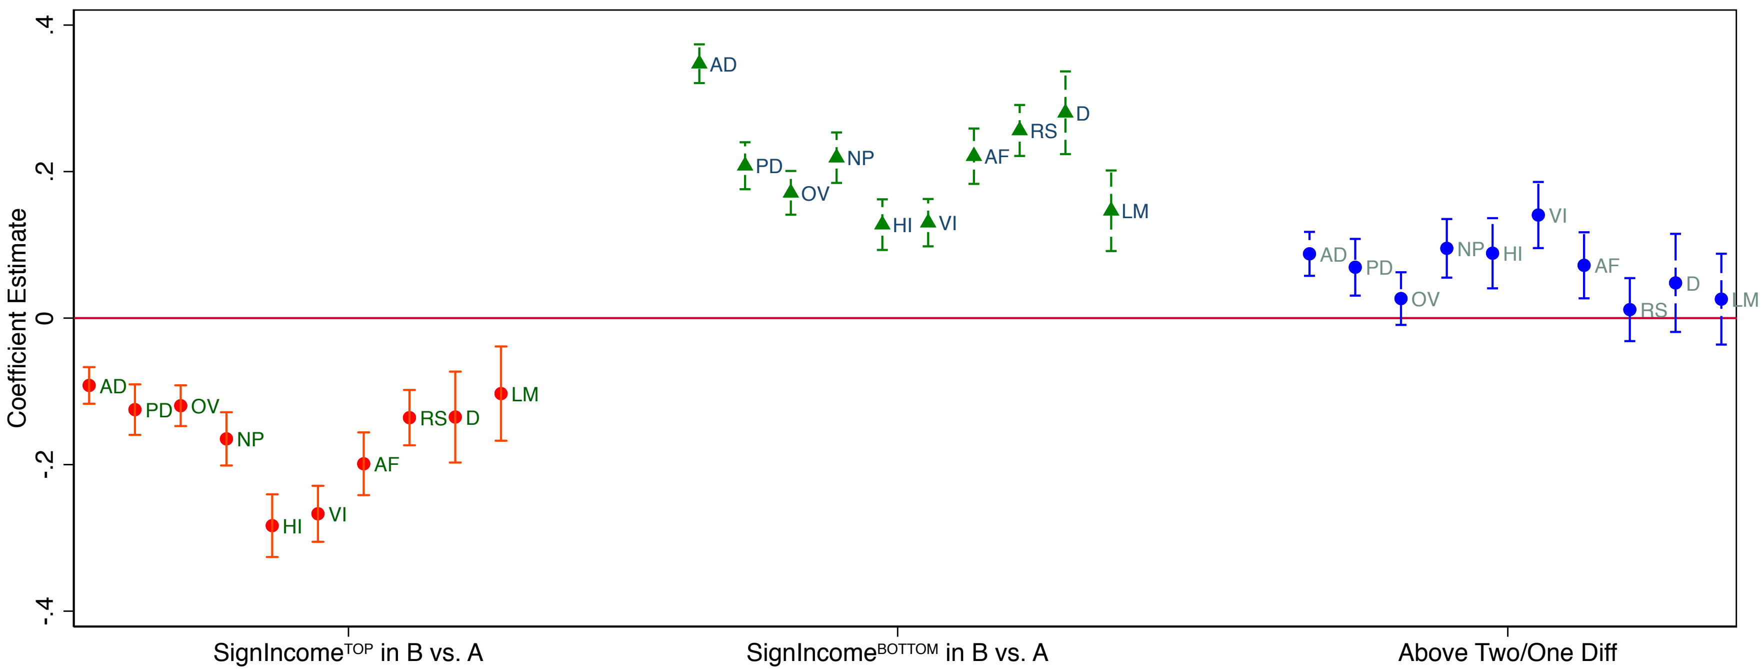

Note: The coefficients plotted reflect those generated by the specification in column (1) of Table 3. The 95-percent confidence intervals plotted are based on standard errors clustered by subject.