Reference points and redistributive preferences: Experimental evidence

Jimmy Charité, Raymond Fisman, Ilyana Kuziemko, Kewei Zhang*

October 3, 2022

Abstract

We explore whether individuals, when acting as social planners, respect others' reference points. We allow subjects to redistribute unequal, unearned initial endowments between two anonymous recipients. Subjects redistribute twenty percent less when recipients know their initial endowments (and thus may have formed corresponding reference points) than when the recipients do not know their initial endowments, in which case we observe near-complete redistribution. The result holds for both within- and between-subject comparisons and is robust to a number of variants in design. The extensive margin response (redistributing zero versus any amount) drives the difference, further suggesting that respect for reference points drives the observed limited redistribution.

Key words: Redistributive preferences, taxation, reference dependence

*We thank Alberto Alesina, Angus Deaton, Stefano DellaVigna, Marc Fleurbaey, Larry Katz, Benjamin Lockwood, Nolan McCarty, David Moss, Howard Rudnick, Stefanie Stantcheva, Matt Weinzierl, and Leeat Yariv for helpful discussions, as well as seminar participants at the AEA meetings, Berkeley, Bocconi, UCSB, Chicago, Columbia, Harvard, LSE, NBER, Northwestern, Princeton, UPF, and Stockholm. Financial support from the Tobin Project is gratefully acknowledged. Adith Srinivasamurthy provided invaluable assistance with JavaScript programming and Dana Scott provided excellent research assistance. Charité: Columbia University (email: jc4144@columbia.edu); Fisman: Boston University (email: rfisman@bu.edu); Kuziemko: Princeton University (email: kuziemko@princeton.edu); Zhang: Renmin University of China (email: zhangkw@bu.edu).

1 Introduction

There exists a large body of evidence that individuals often evaluate their options relative to a *reference point*, often the status quo.[1] It is natural, therefore, to ask whether people show concern for reference points when considering questions of fairness and justice that involve *others*, such as redistribution from higher- to lower-income individuals. Yet relatively limited work on this question exists in the political science or economics literatures.

Introducing the idea of reference dependence to the analysis of redistributive decisions is potentially important. First, if redistributors respect (status quo) reference points, the predictions from textbook models in political economy and public finance can in fact be overturned. As we illustrate in Appendix Figure A.1, with reference-dependent utility functions, because of the large utility consequences of losses relative to gains, it can be optimal to allow substantial inequality to remain, even when endowments are exogenous (and thus moral hazard is not a concern). By contrast, with "well-behaved" continuously differentiable neoclassical utility functions, a utilitarian social-planner would always choose full redistribution in the simplified world of no moral hazard (Mirrlees, 1971).[2]

In this paper, we explore the consequences of reference dependence on the demand for redistribution in an experimental setting. Our main experiment tests whether subjects take into account reference points when asked to play the role of social planner, using a design that reflects a simple two-person redistribution case in which one individual receives a greater endowment than the other. Specifically, we confront subjects (referred to in the paper as "redistributors," though such language is never used in the experiment) with a redistributive

[1]The existence of reference-dependent preferences is a key assumption of Prospect Theory, introduced by Kahneman and Tversky (1979). Since then, countless papers have found reference-dependent behavior in settings as diverse as labor-supply to stock-market transactions. We review the most relevant applications from this voluminous literature later in the introduction.

[2]We are not alone in incorporating non-standard preferences into optimal-taxation analysis. Bernheim and Rangel (2004), O'Donoghue and Rabin (2006) and Allcott *et al.* (2014) consider how optimal *goods* taxation changes when consumers suffer from addiction, lack of self-control, or inattention, respectively. Related work has evaluated the efficacy of behavioral "nudges" to increase *compliance* (e.g., Hallsworth *et al.*, 2017). Engström *et al.* (2015) and Rees-Jones (2017) look at the role of loss aversion in tax compliance.

decision involving two other (randomly selected and anonymous) "recipients" (again, this language is not used in the experiment), who received (based on a coin flip) unequal initial endowments, with one awarded \$5 and the other \$15. We emphasize that the endowments are random and exogenous, thus removing concerns of deservedness or moral hazard, two key factors that could otherwise account for limited redistribution.

To vary redistributors' beliefs about the recipients' reference points, redistributors in the treatment arm were told that the recipients *had already* been informed of their ex-ante allocation; redistributors in the control arm were told that the recipients had *not* been informed of their ex-ante endowments and would thus only be informed of their ex-post allocations. Subjects were then free to redistribute (or not) any whole-dollar amount between recipients, subject to the constraint that all ex-post allocations remain non-negative. We clarify for the subject that neither recipient will ever learn that a third party (i.e., the subject himself) was responsible for the shift in allocation, to minimize the extent to which subjects might be concerned for the inference that recipients might make about the redistributor.[3] Each redistributor was presented with both treatment and control conditions, with the ordering chosen at random, to allow both between-subject (based on the first condition they observe) and within-subject analyses.

Our main finding is that subjects in the reference-point treatment (in which recipients have been informed of their ex-ante endowments) are significantly less likely to reduce inequality between recipients than subjects in the control condition. In our preferred specification, control-group redistributors erase 93 percent of the initial \$10 difference between the recipients' endowments, compared to 77 percent in the treatment group. Importantly, most of this difference is explained by the extensive margin (the decision to redistribute zero versus any other amount), consistent with respect for the ex-ante reference point driving the lack of redistribution in the treatment group.

Our estimates hold across a large number of robustness checks—dropping those who finish

[3]See, for example, Ariely *et al.* (2009), on how social desirability bias in tasks involving concern for others.

the survey in a suspiciously short amount of time, changing presentational aspects of the experiment, and moving from hypothetical to "real stakes" scenarios. While for convenience much of our data are collected via Amazon Mechanical Turk (MTurk), we show that our results replicate on the more representative sample of Americans (by adding our survey to the *Understanding America Study* run by the University of Southern California). We also show that between-subject and within-subject esimates are very similar. Further, the estimates remain stable as we collected data on these variants over the course of several years.

We calibrate the size of our effect relative to what the literature suggests is one of the key barriers to redistribution: merit, the sense that the well-off have earned (and hence deserve) their income.[4] In one session, we ask respondents to redistribute unequal endowments between recipients whose initial allocations are assigned either by a coin toss (control), or earned by correctly answering SAT questions (treatment). In both cases, the recipients would only know their final endowments, so the coin-toss scenario is identical to the control arm of the main experiment. We find that the control group erases 90% of the initial $10 difference and the treatment group erases 57%. Thus, our hypothesized "reference-dependence" mechanism has an effect size that is nearly half of the luck-versus-merit effect, suggesting that it could be an important and heretofore underappreciated explanation for limited redistribution.

There are several candidate explanations for the reference-dependent redistributive decisions of our subjects. Redistributors may (1) view initial allocations as property rights (e.g., Gächter and Riedl, 2005); (2) exhibit a status-quo bias; (3) believe that telling the recipients their initial endowments serves as an implicit promise, and thus reject redistribution on procedural-justice, not reference-point, grounds; (4) respect subjects' reference points (our proposed hypothesis). The near total redistribution in the control condition argues against the most straightforward versions of explanations based on property-rights or status-quo bias. To examine the potential role of procedural justice, we ran a version of the survey which informed subjects that recipients in the treatment condition had been

[4]See, e.g., Alesina and Angeletos, 2005, Durante *et al.*, 2013, and Almås *et al.*, 2020.

told up front that their initial allocations *could change* before the receipt of payment. This phrasing should minimize procedural justice concerns, as redistributors can now reallocate money without worry of breaking an implicit promise to the treatment-group recipients. Our results continue to hold in this variation, indicating that procedural justice is unlikely to be the primary explanation for our treatment effect. We conclude that respect for others' reference points is the theory that best explains our full collection of results.

The main goal of our paper is to establish the importance of reference points for redistributive decisions in an experimental setting stripped of as many competing factors as possible, even though such a setting is by necessity quite removed from a real-world scenario. We close the paper with a vignette experiment that aims to connect our results from the lab to a more policy-relevant and realistic setting, though we emphasize that these latter results are only suggestive and we hope may serve to encourage future work. The vignette describes a person who received a substantial raise via a lucky event – as a result of a merger between her employer and a larger, better-paying firm. Subjects in the control group are told that the worker *just* received the raise this year and subjects in the treatment group are informed that she received the raise five years ago. Both groups are then asked how much this person should pay as an average income tax rate. Subjects confronted with the treatment vignette set lower tax rates on average, which can potentially be interpreted as their respecting the more deeply embedded, five-year reference point of higher post-tax income.

Our paper contributes to a large literature on reference points and the endowment effect, but we believe its application to redistribution from the rich to the poor is quite understudied. For the most part, papers on the endowment effect ask the subject to interact with a single "other" in a dictator-game setting, whereas we put the subject in the role of a Mirrleesian social planner who is tasked with making redistributive decisions between ex-ante "richer" and "poorer" recipients. Thus, we set up an experiment in which subjects are confronted with two competing motives, both of which have been shown to have strong effects on behavior in past experiments: inequality-aversion (which pushes toward redistribution) and respect for reference points (which pushes against it), removing any consideration of own-payoffs, which complicates the interpretation of standard dictator games.[5]

As such, our paper contributes to the political-economy and public-finance literature that aims to understand why demand for redistribution falls short of the levels predicted by standard models, a literature that political scientists and economists have both advanced. In the standard optimal taxation framework (Mirrlees, 1971), the social planner wants to redistribute because the marginal utilities of the poor are greater than those of the rich, but is deterred from full redistribution by supply elasticities (the "supply" of labor or taxable income declines with the tax rate). However, recent work suggests that standard estimates of labor-supply and taxable-income elasticities with respect to net-of-tax rates yield higher optimal marginal tax rates than are typically observed.[6] If voters are altruistic or are difference-averse—and many lab experiments suggest subjects indeed are—the link between inequality and demand for redistribution is expected to be even stronger (see, e.g., Dimick *et al.*, 2018 and Lü and Scheve, 2016).

Moving from the normative to the positive, a democratic political process would similarly be expected to generate a high degree of redistribution: given a right-skewed pre-tax income distribution, the majority of voters have an incentive to support high effective tax rates on the wealthy (see Meltzer and Richard, 1981, and Dimick *et al.*, 2016 for how support may vary by own income). Past work aiming to explain this reluctance has focused on the prospect of upward mobility (Benabou and Ok, 2001), the effects of "policy-bundling" redistribution with other, cross-cutting issues (typically race in the U.S. context, see Lee and Roemer, 2006), and the public's misinformation about income inequality (see Ariely and Norton, 2011 on the level of misinformation, though see also Kuziemko *et al.* (2015) on the limited effect of information on policy preferences).[7]

[5]See List (2007) for a review of dictator games and their sensitivity to the choice set offered subjects.

[6]See e.g., Diamond and Saez (2011), who argue that a utilitarian social welfare function would yield top marginal tax rates over 70 percent given consensus estimates of labor supply elasticity.

[7]For one closely related contribution, see Almås *et al.* (2020), who, like us, put subjects in the

We provide a heretofore underexplored explanation for the limited demand for redistribution that, our results suggest, may be quantitatively important. Our experimental findings relate to prior theoretical work which shows that optimal income tax results may change when agents are loss averse (see, e.g., Kanbur *et al.*, 2008; Dai, 2011). Our results offer evidence that the limited-redistribution-of-exogenous-endowments result can be rationalized in an optimal-tax model with a standard utilitarian social-welfare function, albeit with non-standard individual utility functions. Further, there is also a somewhat older literature that examined optimal tax results when agents care about their income [8] Our paper is distinct from existing work, by combining our experimental approach with the focus on the implications of "behavioral" agents on decisions, such as those over redistribution, that impact others.

Given that our explanation for limited redistribution builds on insights from social psychology, our paper also contributes to the field of behavioral public finance. One strand of this literature has looked at the implications of "behavioral agents" for the taxation of *goods*.[9] A distinct literature has considered the implications of behavioral models for tax compliance.[10] Recently, Jones (2020) has shown that a homeowner is more likely to contest a property-value assessment if it is above the reference point that he argues is created by the previous year's assessment, a particularly convincing example of loss-aversion in tax compliance because contesting assessments is an effort-intensive, real-world task.

role of social planner, but find that the starting points of the two recipients do not appear to deter redistribution. As we argue in Appendix C, where we provide a more detailed comparison of our paper and theirs, recipients in Almås *et al.* (2020) are explicitly told that these starting points are subject to change, thus we suspect the social planner feels little need to protect the richer party from loss aversion.

[8]There is, in addition, a somewhat older literature that derived optimal tax results when agents care about own income relative to others (see, e.g., Boskin and Sheshinski, 1978; Oswald, 1983).

[9]Past work has examined the implications of addiction, lack of self-control, and inattention. See, e.g., Bernheim and Rangel (2004), O'Donoghue and Rabin (2006) and Allcott *et al.* (2014), respectively.

[10]A number of studies consider the efficacy of "nudges" to increase compliance (e.g., Hallsworth *et al.* (2017)); most directly related to our work, Engström *et al.* (2015) and Rees-Jones (2017) look at the role of loss aversion in tax compliance.

The paper proceeds as follows. The next section describes the design of the main experiment. Section 3 provides details on the data-collection process. Section 4 presents the main results, along with robustness tests and related results. Section 5 describes a companion experiment on preferred tax rates. Section 6 highlights the challenges that our results pose to the inclusion of reference points in social-welfare analysis. Section 7 concludes and offers ideas for future work.

2 Experimental design

We collect data from eleven distinct sessions (all described in Appendix Table A.1). While there were some small differences in wording and presentation among these sessions (which we introduced to test robustness and we describe later in the paper), in each session subjects are asked to allocate a fixed surplus between two other subjects. Five sessions also include our tax-survey question (described in Section 5). One session involved a luck-versus-merit rather than the reference-point treatment (described in Section 4).

Each experimental session presents subjects with the opportunity, in most cases hypothetical, to transfer money between two other anonymous participants. In all cases, *the redistributor received only his show-up fee regardless of his decision, so he has no direct self-interested motivation.*[11]

Respondents randomized into the control arm of the survey encountered the following instructions:

> Consider two other participants on MTurk, person A and person B. Based on a coin flip, we have given $5 to person A and $15 to person B.
>
> You can now transfer money between persons A and B. *Persons A and B are not told how much money they were initially given*. If you decide to give Person A $X instead of $5, he or she will simply be told that they have been given $X, and will not know how much they started with. Nor will they know that there is another person (Person B) involved, or that a third party (you) determined the money they received.

[11] Our approach is thus similar to the "spectator design" used in Almås *et al.* (2020).

> Please indicate below what transfer, if any, you would make.

A slider and interactive bar graph (which reflects in real time movements of the slider) appeared directly below these instructions, allowing respondents to easily and transparently transfer money between players. The default position of the slider was on the ex-ante (\$5, \$15) distribution. Appendix Figure A.2 provides a screenshot.

For those randomized into the treatment arm, the second paragraph of the control instructions was modified as follows:

> You can now transfer money between persons A and B. *Persons A and B have already been told how much money we have given them*. If you decide to give Person A \$X instead of \$5, they will be told that they now have \$X instead of \$5. They will not know that there is another person (Person B) involved, or that a third party (you) determined the money they received.

Appendix Figure A.3 provides a screenshot.

To test the robustness of the results to within-person instead of between-person variation, we also performed the "reverse experiment," and so immediately after answering the treatment (control) version of the question, treatment (control) respondents answer the control (treatment) version of the question (with the labels "Persons A and B" replaced with "Persons C and D").

A version of the survey where the redistribution experiment is presented first is available in Appendix E.

We emphasize in both treatment and control conditions that the recipients will never know that a third party (the subject) determined the final allocations. We wanted to separate our reference-point effect from any effect driven by subjects' concerns over recipients forming a negative impression of the redistributor, as there is a large literature suggesting that subjects behave differently if other subjects (or an audience more generally) can observe their behavior. While this body of work generally focuses on choices between fairness and self-interest (and in our setting, self-interest plays no obvious role, as the subject receives only her show-up fee regardless of how she splits the surplus), we nonetheless wanted subjects to make their decisions with the assurance of anonymity.[12] In Section 4.6 we explore the open-ended answers respondents gave in explaining their thinking, and none mention concern about what the recipients thought of them, suggesting that our attempts to assure anonymity were successful.

Two additional presentational aspects of the main experiment deserve mention. First, the use of the slider requires a default position, which we set to the status-quo allocation of $5 and $15 dollars. As such, we suspect that anchoring bias could lower the amount of redistribution in both treatment and control and thus attenuate any treatment effect. Second, to illustrate clearly the treatment scenario, we write: "If you decide to give Person A $X instead of $5, they will be told that they now have $X instead of $5." By using the poorer person as the illustration, if anything we should prime redistributors to think of the *pleasant surprise* that the person starting with $5 will experience, again biasing the experiment *against* finding our hypothesized effect.

The survey ended with standard demographic questions, as well as a question on preferred candidate in the most recent presidential election (the 2012 election for the MTurk subjects, and the 2016 election for the *UAS* subjects). These questions allow us to examine whether our treatment effect is larger for certain groups, and also to compare our sample to more representative populations such as the General Social Survey.

The final questions of the survey relate to whether respondents felt any part of the survey was confusing or biased and also asked for any other feedback they wished to share.

[12]See, e.g., Andreoni and Bernheim (2009), Dana *et al.* (2007), and Ariely *et al.* (2009), and papers cited therein. Even if subjects do not trust our assurance of anonymity, or skip through it when they skim the instructions, it would still seem to suggest a concern for reference points, as they must implicitly weight more the (negative) judgment of the person receiving less than their expected $15 than the (positive) judgment of the person receiving more than their expected $5. Whether our reference point effect is *enhanced* by weakening the promise of anonymity is an interesting question for future work.

3 Data

We recruited and compensated most of our subjects through Amazon's Mechanical Turk (MTurk) marketplace (which we describe in detail in Section 3), but redirect them to surveys that we built with Qualtrics' online survey software, adding functionality with JavaScript as needed. Our MTurk data was collected in 2014. Cognizant of concerns that MTurk is not representative of the U.S. adult population, in 2019 we replicated our analysis by adding questions to the nationally representative *Understanding America Study* run by the University of Southern California.[13]

MTurk is an online labor market where "requesters" can post *human intelligence tasks* (HITs) to be completed by "workers." Social scientists have increasingly used MTurk to perform experiments and collect survey data (see Kuziemko *et al.*, 2015 and citations therein for a review), and as such we relegate most of the details of our data-collection procedure to Appendix D. In the Appendix, we describe the steps we take to ensure a subject pool composed of attentive adult Americans (actual humans as opposed to "bots") taking the survey in good faith. As we discuss in the Appendix, beginning in 2018 MTurk saw a surge of bots posing as real humans, but bots were much less of a problem in 2014 when we collect our MTurk data.

While we discuss representativeness and experimental balance in more detail later, we note here that the data pass basic reality checks (for example, subjects that report having supported Mitt Romney in 2012 tend to be white and male, mirroring patterns observed in polling data). Over three-quarters of respondents went on to answer an open-ended "feedback" question, with the vast majority providing positive feedback on the survey and writing in colloquial, American English.[14]

[13]In the UAS survey, subjects only completed the experiment after answering questions on wealth and income taxation (see Fisman *et al.*, 2020). Thus, as we discuss later, the smaller effect sizes we observe in this sample could result in part from respondent fatigue. As seen in Table A.1, the main experiment also was preceded by the tax vignette experiment in three instances, which also tended to have smaller effect sizes.

[14]We suspect that the positive feedback likely reflects the tedium of most other mTurk tasks. As

Two important concerns in our setting are experimenter-demand and social-desirability effects. Thus, it is important to assess whether subjects believed that we, the researchers, wanted to find a result that favored a particular political slant. In Appendix Table A.2 we show how questions on perceived political bias of the survey vary with treatment status. The vast majority of respondents (85 percent) report that the survey felt unbiased to them, with ten percent detecting left-wing bias and four percent right-wing bias.[15] Important for the interpretation of our findings, these reports of perceived bias are uncorrelated with treatment status. Similarly, survey fatigue should not affect our estimates of the treatment effect, as we show in Appendix Table A.2 that the average number of minutes taken to complete the survey is also independent of treatment status.

While we believe that our MTurk data are of high quality in that real human subjects answered the questions seriously and in good faith, we also replicate our results on the more representative data collected through the *Understanding America Study* (UAS), run by the University of Southern California.

Appendix Table A.3 provides details on the subjects who completed our survey experiments, separating between MTurk and *UAS* subjects and comparing them to the nationally representative sample of adults in the General Social Survey. Consistent with past work using MTurk, younger and male subjects are over-represented, and despite being more educated than the representative American adult, subjects have lower household incomes. On these dimensions, the *UAS* data are more representative, and are quite close to the GSS averages.

Appendix Table A.4 provides a larger number of covariates for MTurk and UAS participants (not limited to those that can be compared to the GSS) and reports differences between the control and treatment groups. There is good experimental balance, with no variable showing a statistically significant difference at the five-percent level. In particular, variables we suspect could impact redistributive decisions in our setting—political prefer-

noted, essentially all of the negative feedback concerned the difficulty of some of the "captchas."

[15] There was an option to select "bias of some other type," which a handful of subjects chose.

ences as captured in presidential election votes—are very similar across experimental arms.

4 Results from the redistribution experiment

4.1 Main results

Table 1 shows, for the full sample across all ten sessions of the redistribution experiment (including both the MTurk and UAS samples), the main between-subject differences in total redistribution for those first assigned to the control versus those first assigned to the treatment. Recall that redistributing $5 from the "richer" to "poorer" recipient would result in strict equality ($10, $10). Column (1) shows the treatment effect controlling only for session fixed effects. Those in the control group achieve nearly complete redistribution, shifting an average of $4.35 from the richer recipient to the poorer one, or 87% of the redistribution required for strict equality. Recall that the default position of the slider was the status quo ($5 and $15) allocations, suggesting that anchoring or inattention would bias the control group results *against* inequality-reducing redistribution. Those assigned to the treatment redistribute on average $0.75 (or 17 percent) less than those in the control.

Column (2) drops subjects who finished the survey in less time than one could reasonably be capable of completing it.[16] We view this result as our preferred specification: redistribution is 18 percent lower in the treatment than in the control group. Column (3) further excludes subjects for whom the $5/$15 experiment was not the first survey item survey (removing subjects that may be contaminated by exposure to our income tax survey experiment, discussed in Section 5).

We show our main results separately by session in Panel (a) of Figure 1. Sessions 2 and 3 are pure replications of Session 1. For Sessions 4 and 5, the $5/$15 experiment is identical to that described in Section 2, but it appeared *after* our income-tax experiment (described in the next section).

[16]Less than three minutes for the first session, which had fewer follow-up questions, and six minutes for other sessions.

Table 1: Chosen redistribution in the \$5/\$15 reference-point experiment

	Amount redistributed			Zero Redistribution		
	(1)	(2)	(3)	(4)	(5)	(6)
Treated in first stage	-0.747***	-0.789***	-0.829***	0.136***	0.134***	0.158***
	(0.100)	(0.103)	(0.134)	(0.015)	(0.016)	(0.019)
Cont. gp. mean	4.345	4.374	4.682	0.119	0.118	0.061
Ex. short duration	No	Yes	Yes	No	Yes	Yes
Ex. presented second	No	No	Yes	No	No	Yes
Observations	2352	2194	1227	2352	2194	1227

Notes: The outcome in the first three columns is the amount of dollars redistributed from the recipient who starts with \$15 to the recipient with \$5, and the (binary) outcome in the final three columns is whether the subject chose zero redistribution. "Treated" refers to the subject being told that the recipients knew their initial endowments (as opposed to the control group who were told that the recipients would only know their final allocations). All regressions include session fixed effects. Ex. short duration: exclude subjects who finish the survey in a suspiciously short amount of time. Ex. presented second: exclude survey sessions where the main \$5/\$15 reference-point experiment was not presented first. $^{*}p < 0.1,^{**} p < 0.05,^{***} p < 0.01$.

The remaining sessions test the sensitivity of our result to changes in wording, presentation and sampling platform. While in all other cases our subjects dealt with hypothetical situations, in Session 6 they were informed that with ten-percent probability their decisions would be implemented on actual individuals. The main effect remains significant at the five-percent level.[17] In another robustness check (Session 7), none of the text was italicized or underlined, and the underlined reminder message (see Appendix screenshots) placed next to the slider was removed. This "no emphasis" version is significant at the ten-percent level ($p = 0.071$).

In Session 8, we altered the language in the treatment condition to convey to redistrib-

[17] We cannot know for certain whether our subjects realized that the experiment involved hypothetical individuals, though in the experiment the decision is worded in the conditional ("Please indicate below what transfer, if any, you would make"). Moreover, in the MTurk invitation, subjects were told it was an "opinion" survey, again suggesting hypothetical scenarios. Note that in Session 6, the round involving real stakes, we do not employ deception. We indeed give money to MTurk workers in a manner consistent with the decisions of our subjects. In the case of the treatment condition, recipient MTurk workers were told initially they had \$15 (or \$5) and subsequently told that instead they had whatever amount our subjects allocated them. In the case of the control condition, they are just given sums of money.

utors that the initial endowments *should not be seen as a promise or obligation* to persons A or B. As Leventhal (1980) notes, the procedural justice "rule of adhering to commitments...dictates that fairness is violated unless persons receive that which has been promised to them." To limit redistributors' sense of commitment to an initially promised allocation, we modified the wording of the underlined portion of the instructions to read: "Persons A and B were told how much money they were initially given, though they have also been told that the amount might increase or decrease." All other text was unchanged. As shown in Figure 1, this "Procedural Justice" version of the experiment has a somewhat smaller treatment effect than the average, but remains significant at the ten-percent level ($p = 0.096$). Given that we attempt to soften the reference point somewhat in this version, we view the the survival of a reference-point effect (albeit at the ten-percent level of significance) to be an indication that our subjects respect reference points even when no explicit promise is given to the recipients.

Finally, to ensure that results are not driven by some peculiarity of the MTurk platform, we added questions to the *Undestanding America Survey*, a panel comprised of a representative sample of Americans, and managed by researchers at USC (Session 10). The treatment effect is similar to those of the other sessions, but with a p-value of 0.114.

The stability of estimates in Figure 1(a) highlights the robustness of our results to various changes in presentation and different samples. While the estimated treatment effects differ somewhat across sessions, we cannot reject equality of any pair of estimates at even the 10% level.

Of course, Figure 1(a) also makes clear that not every session produces a result that is significant at the five-percent level. Four of the ten sessions yield a treatment effect that is not significant at the five-percent level, and two fail to be significant at the ten-percent level. In Appendix B we investigate the expected number of "failed" replications (with 95-percent and 90-percent confidence as the standard) under the assumption that our pooled treatment effect represents the true data-generating process. We simulate data-collection processes to

Figure 1: Treatment effects for the reference-point experiment, by session

(a) Total amount redistributed to ex-ante poorer subject

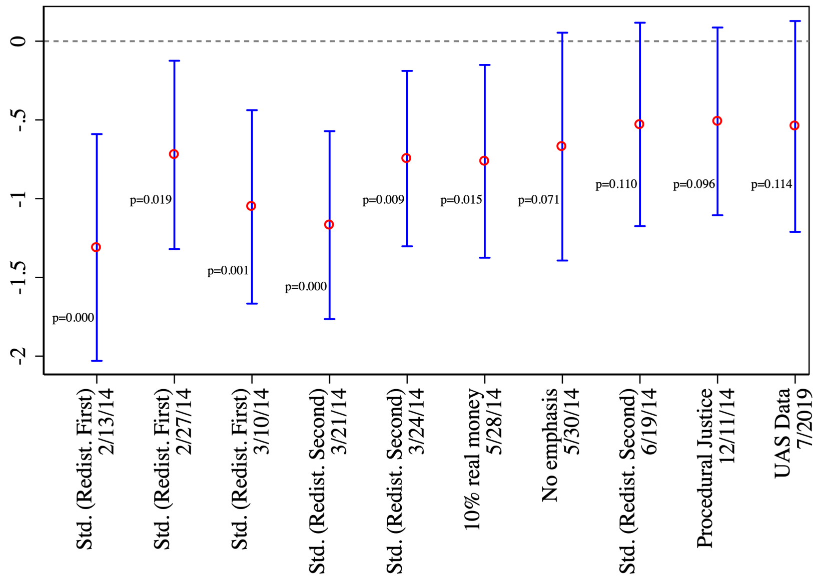

(b) Zero redistributed (dummy variable)

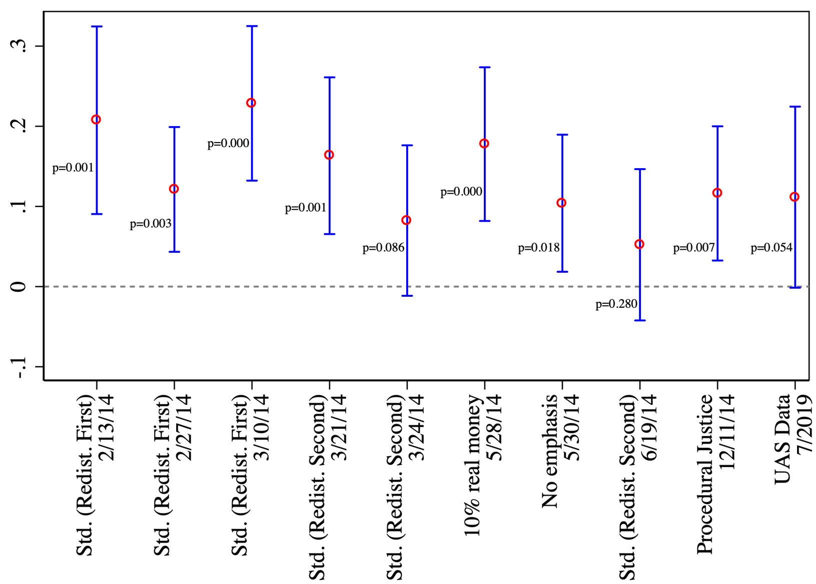

Notes: This figure shows treatment effects and ninety-five-percent confidence intervals, separately for each of the ten sessions of the reference-point experiment (using the sample restrictions in col. 2 of Table 1). We show the treatment effects for two outcomes: (a) the amount redistributed and (b) the (binary) outcome of choosing to redistribute $0. The first five sessions of the experiment have the exact presentation as outlined in Section 2. The next four sessions are variants (described in Section 4) included to probe robustness. The final session has the standard presentation, but subjects are taken from the more representative Understanding America Study instead of MTurk. In all sessions, we drop subjects who finish the survey too quickly.

reflect ten rounds of data collection with sample sizes that match those in our experiments. We find that on average 3.47 rounds will not produce a treatment effect significant at the five-percent level and 2.45 rounds will not produce a treatment effect at the ten-percent level.

We provide greater detail in Appendix B, but in summary the simulations produce results that are very close to what we observe in our actual data. We note as well that our simulation assumption—that all sessions should have the same treatment effects—is conservative, as we in fact have sessions that include refinements we would *ex ante* expect to weaken the treatment effect (in particular, the "No Emphasis" and the "Procedural Justice" versions).

4.2 Intensive v. extensive margin effects

Figure 2 shows histograms of the final allocation for the ex-ante "poorer" player, for the treatment and control groups. For both groups, the distribution is bimodal, with the mass at (10, 10) but also a second, shorter peak at (5,15); there is almost no mass between these two points. Thus, most of the treatment effect occurs at the *extensive* margin, i.e., the decision to redistribute at all. The lack of intermediate choices is broadly consistent with subjects responding to recipients' reference points, in that partial redistributions *both* fail to equalize endowments (and thus do not satisfy the desire to avoid inequality) *and* lead to disappointment for the "richer" individual with the $15 reference point.

The final three columns of Table 1 shows these extensive-margin results in a regression framework, in which the dependent variable is an indicator denoting the subject chose no redistribution at all. Just under twelve percent of control-group respondents choose to reallocate zero dollars, compared to over 25 percent (0.119 + 0.136) in the treatment group, a highly significant difference. The remaining columns of Table 1 show that this result passes the same robustness checks as the main result.

Panel (b) of Figure 1 shows that the extensive-margin result is stable across all ten sessions. We note that the replication fails at the five-percent level of significance in three individual rounds (and one fails at the ten-percent level); we again show in Appendix B that

Figure 2: Histogram of ex-post allocations for the ex-ante poorer subject

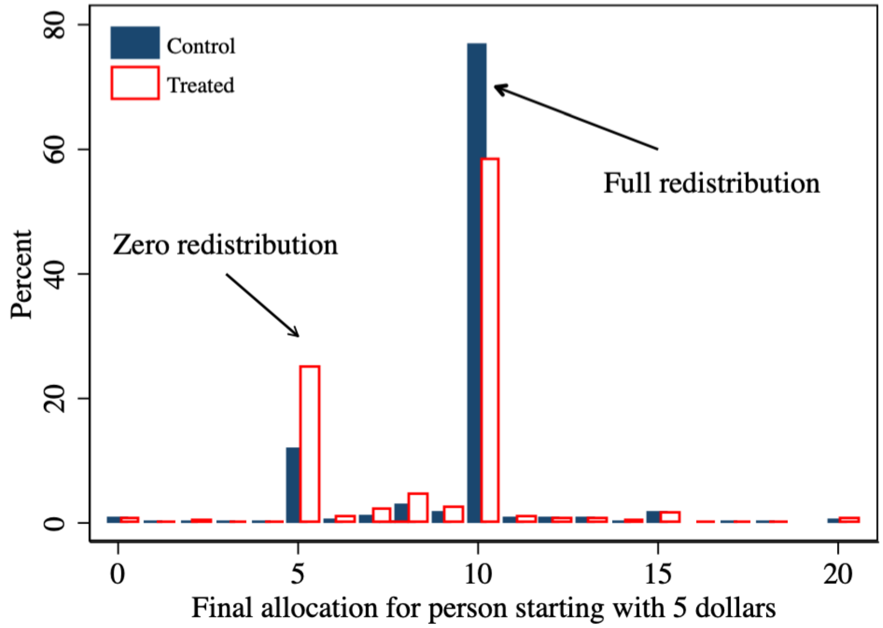

Notes: Sample used in the figure is that in Column 2, Table 1. Treated refers to subjects who were told the recipients knew their initial allocations and control refers to subjects who were told that the recipients would only know the final allocation.

this "failure rate" is consistent with by-round sampling variation if the true data-generating process is captured by the pooled treatment effect in Table 1.

4.3 Results from within-subject variation

As noted in Section 2, we presented subjects with the reverse condition immediately after their response to their initial scenario (i.e., those assigned to the control condition were then confronted with the treatment scenario, and vice-versa). Figure 3 uses this within-subject variation and shows the main results as well as the extensive-margin results separately by session (with standard errors clustered by subject).

In the main specification, all sessions produce negative treatment effects significant at the ten percent level, with only the "Procedural Justice" variant failing at the five-percent level ($p = 0.089$). In the extensive-margin specification, all ten rounds produce negative treatment effects that are statistically significant at the five-percent level.

While we generally focus on the between-subject analysis, we find the robustness of the within-person results to be very reassuring. In Appendix Table A.5 we show that the pooled within-subject treatment effect is quite similar to its between-subject analogue (0.661 versus 0.789) and is similar for those who first confront the treatment condition and then the control (0.60) and for those who first confront control then treatment (0.71).

4.4 Related results and robustness

In both lab experiments and surveys, individuals' desire to redistribute is more muted when they believe that the subjects in question have *earned* their incomes via merit or effort rather than gaining their incomes via luck or chance.[18] To get a sense of the magnitude of our reference-point effect, we compare our main results to those from an almost identical experiment in which, instead of varying the salience of reference points, we vary whether luck or merit determined the initial endowments of Persons A and B.

To make this comparison, we run a session in which the control arm was kept the same (i.e., the \$5 and \$15 endowments were determined by a coin flip, and the recipients know *only* their final allocations), while in the treatment arm respondents were told "the initial amounts given to Persons A and B *were based on their performance on SAT questions* [emph. in original]." Importantly, as with the control version, the redistributor is told that Persons A and B would only learn of their own final allocation (i.e., we shut down any reference-point effect in both the treatment and control arms of this particular round).

In Table 2, we compare the results from this round to those in our standard reference-point experiments. As we would expect given past work, perceived merit (as proxied by SAT performance versus luck) is a large deterrent against redistribution. Those who were told that Persons A and B received their endowments due to SAT performance redistribute \$1.81 less than the control group. Thus, our reference-point effect (reproduced in column 2)

[18]See, e.g., Alesina and Angeletos (2005), Durante *et al.* (2014), Barber IV and English (2019) and citations therein.

Figure 3: Within-subject treatment effects for the \$5/\$15 reference-point experiment

(a) Total amount redistributed to ex-ante poorer subject

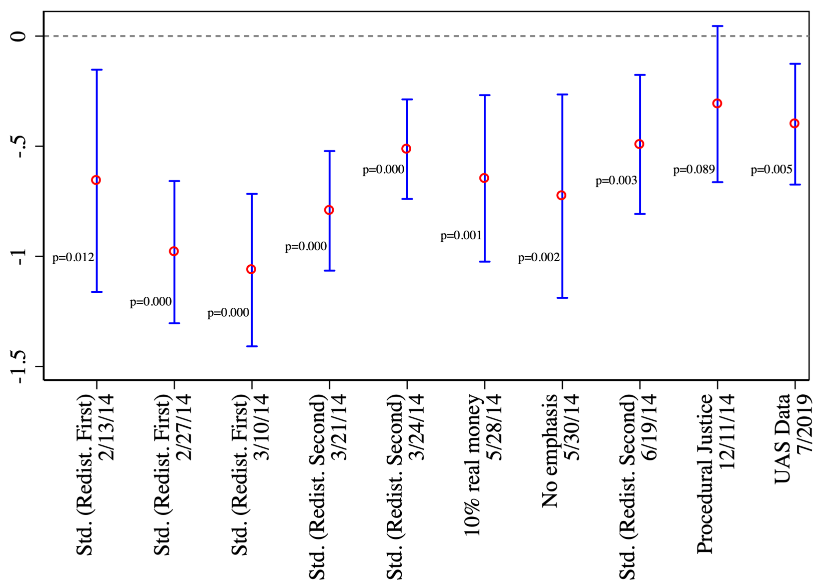

(b) Zero redistributed (dummy variable)

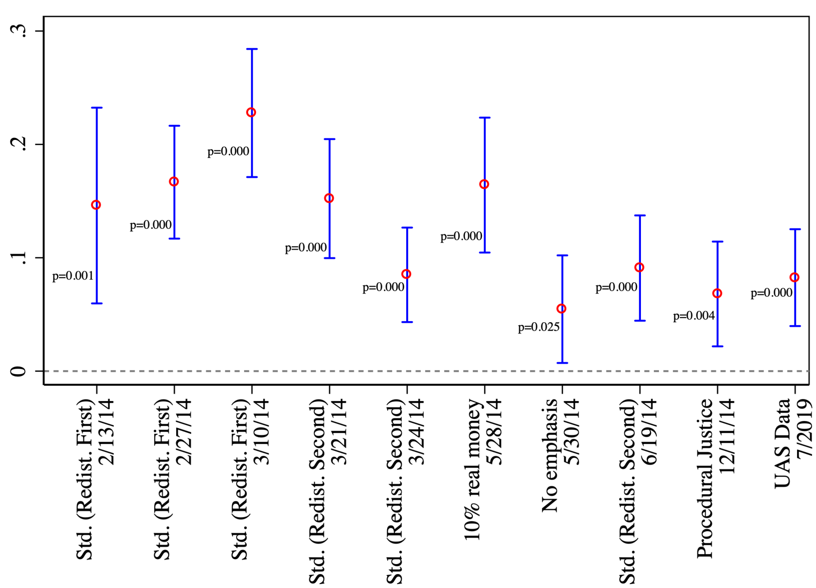

Notes: This figure replicates the analysis from Figure 1, but uses within-subject variation instead of between-subject variation in treatment status. Ninety-five-percent confidence-intervals based on standard errors clustered by subject.

is 42% the size of the merit effect.

We noted in the previous subsection that sensitivity to reference points should lead to a large *extensive-margin* response (i.e., full redistribution or no redistribution), as in-between allocations fail both to achieve equality and to avoid disappointment for the recipient who receives less than she expected. In the SAT-versus-coin experiment, in-between allocations will not disappoint the person who initially has $15, since recipients only learn their final allocations in both the treatment and control arms of this experiment.

The remaining columns of Table 2 test whether the extensive margin is indeed less important in explaining limited redistribution in the SAT experiment than in the main experiment. Column (3) shows that for the SAT experiment, adding the extensive-margin dummy variable of zero redistribution reduces the magnitude of the coefficient on the treatment dummy, but more than fifty percent of the effect remains, and the coefficient on $Treated$ remains highly significant. Column (4) shows a very different pattern for the main reference-point experiment: adding the zero-redistribution dummy explains all but 15 percent of the original treatment effect, and the coefficient on $Treat$ is now only marginally significant despite a much larger sample than in column (3). Column (5) shows that the differences between cols. (3) and (4) in how much more of the treatment effect remains after accounting for the extensive-margin dummy is indeed statistically significant. After accounting for extensive-margin effects, in the SAT experiment redistributors still give $.81 less to the $5 recipient than in the reference-point experiment and this difference is significant at the 1% level.

In summary, the SAT experiment helps us gauge the magnitude of our reference point effect—it is not quite half the size of what is perhaps the most robust and well-documented reason individuals shun full redistribution. Furthermore, comparing the patterns of intensive- versus extensive-margin effects helps bolster the argument that respect for others' *reference points*, and not some other mechanism, explains our main effects. When the recipients know their original endowments (as in the treatment arm of our main experiment) and thus reference points can conceivably play a role, respect for the exact starting point of the

Table 2: Comparing main reference-point results to luck-versus-merit experiment

	Dept. var: Amount transfered to person starting with \$5				
	(1)	(2)	(3)	(4)	(5)
Treated in first stage	-1.806***	-0.789***	-0.941***	-0.140**	-0.140**
	(0.360)	(0.103)	(0.260)	(0.070)	(0.071)
Redistributed zero			-4.453***	-4.829***	-4.829***
			(0.307)	(0.093)	(0.094)
Treated × SAT round					-0.802***
					(0.243)
Redistributed zero × SAT round					0.376
					(0.291)
Cont. gp. mean	4.515	4.374	4.515	4.374	4.385
Rounds included	SAT/coin	Ref pt.	SAT/coin	Ref pt.	All
Observations	206	2194	206	2194	2400

Notes: "Treated" in the SAT rounds refers to the subject being told that the recipients were assigned their initial endowments based on performance on SAT-type questions. "Treated" in the reference-point experiment refers to the standard treatment (being told that the recipients know their initial endowments). In all experiments, the control group is told that the recipients do not know their initial endowments and that the endowments were determined by a coin flip. "Redistributed zero" is a dummy for choosing to transfer zero between Persons A and B and is included to separate intensive- and extensive-margin effects. "SAT round" is a dummy variable for the subject being observed in the SAT-versus-coin-toss round. In all columns, those who took an earlier version of the experiment and those who finish in an unrealistically short amount of time are dropped. All regressions that include more than one session include session fixed effects (cols. 2, 4 and 5). $^{*}p < 0.1, ^{**}p < 0.05, ^{***}p < 0.01$.

higher-paid recipient explains a larger share of subjects' redistributive decisions than when recipients do not know how they started (as in the SAT-versus-luck round).

4.5 Heterogeneity

Finally, Figure 4 examines whether the magnitude of the treatment effect depends on demographic and background characteristics. Despite the fact that we accumulated a relatively large sample size over our ten sessions, none of the various dimensions of heterogeneity we consider produce statistically significant differences in treatment size—men and women, whites and non-whites, those above and below our median age or median income, those with and without a college degree, Democrats versus others. Across all of these splits of the data, subjects react to our experiment in a comparable manner.[19]

We emphasize that many of these attributes (especially gender and party identification) have been shown in numerous other experimental and survey-based studies to be highly predictive of tolerance for inequality (e.g., compared to men, women are more willing to sacrifice overall surplus for greater equality; see Fisman *et al.*, 2017). The *lack* of heterogeneity we find suggests that different models of fairness motivate respect for reference points than motivate the traditional equity-efficiency trade-off.

4.6 Discussion of alternative mechanisms

We find evidence that in their role as social planner, subjects' decisions are affected by whether recipients are aware of their initial endowments. There are several primary candidate explanations for this treatment effect. Redistributors may (1) respect property rights over initial endowments, (2) exhibit standard status-quo bias, (3) view initial endowments as a commitment or promise to recipients that they do not wish to break, (4) consider the "consumption commitments" of the recipients and (5) assume recipients have reference-dependent preferences. The data are harder to reconcile with either of the first two explanations. The most straightforward property rights explanation is inconsistent with the near-complete redistribution in our control condition. This result also casts doubt on pure status quo bias,

[19] Beyond demographic variables, a natural source of heterogeneity is subjects' *own* loss aversion. We included some questions in our survey which aimed to capture own-loss-aversion. We attempted to do so in a way that avoided the cognitively taxing and time-consuming approaches of Chapman *et al.* (2018) (or its less efficient predecessor, Abdellaoui *et al.* (2007)), focusing primarily on comparing subjects' responses to simple vignettes that frame situations in terms of losses versus gains. We do not find that individuals who exhibit stronger own-loss-aversion have a stronger treatment effect in our main experiment. However, we also find that measures of own-loss-aversion are relatively uncorrelated with one another, raising questions about what underlying attribute or preference they are capturing. A description of these results are available in an earlier working paper version of this paper, Charité *et al.* (2015). Since measuring own-loss-aversion is an active area of research, we hope that as better measures are developed, and/or populations with well-measured loss aversion become available to the research community, we may be able to revisit this question in the future.

Figure 4: Differential treatment effects from between-subject results

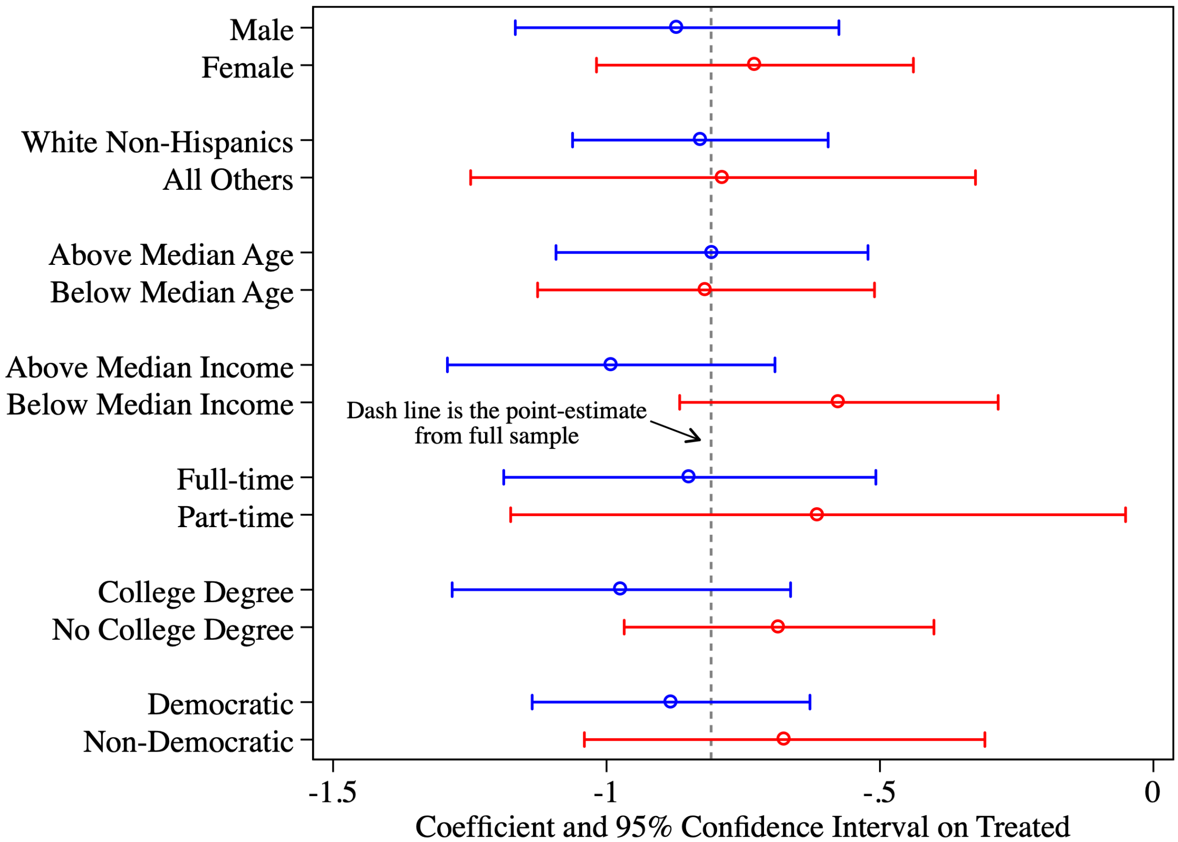

Notes: This figure shows treatment effects and 95% confidence intervals, separately for each of the subgroup of the full sample. All regressions include session fixed effects and the following covariates: age, female, white, black, Hispanic, Asian, income, student status, full-time status, part-time status, Democratic-candidate support in the most recent presidential election, and college degree. Subjects who finished the survey very quickly were excluded from these regressions.

which would also limit redistribution in the control scenario.[20]

Of course, applying the notion of "property rights" is complicated in our control condition: can a recipient have "property rights" over claims that they are unaware that they possessed? In the "Procedural Justice" variant of the experiment (Session 8), however, "property rights" in the treatment condition are deliberately weakened even further by warning the recipients that their ex-ante allocations could change. Yet, while somewhat smaller, the treatment effect remains significant at the ten-percent level, suggesting that a respect for even a broad sense of property rights is not the main constraint on redistributors in the treatment condition. Moreover, this "no promises" variant also argues against a dominant role of distributive obligations as a result of commitments or promises.

Chetty and Szeidl (2007) present a model of consumption commitments that could similarly diminish redistribution by a social planner who takes into account the commitments of relatively well-off individuals. As they observe, however, in a context such as ours it is implausible that *actual* consumption commitments could drive subjects' decisions—the individuals over whom they were making decisions were given money that, by construction, had not yet been spent.[21]

Thus, we argue that respect for others' reference-dependence most easily reconciles our full set of results. This interpretation is further reinforced by subjects' own descriptions of how they made their choices, based on the following open-ended question at the end of the survey: "For the questions where you had to transfer money between other people, how did it change your thinking if these people had already been told how much money they were initially given?"

We focus on subjects who chose different allocations in the two versions of the exper-

[20]See Trump (2015) for experimental work on a particular form of status quo bias—how respondents might deem high levels of inequality legitimate because of a belief in a "just world."

[21]If individuals develop reference-dependent heuristics as a result of commonly observing consumption commitments in their day-to-day lives, it could help to provide an underlying model for reference dependent preferences. Examining this possibility may be an interesting direction to pursue but is outside the scope of our paper.

iments, since they are most relevant for understanding how subjects' thinking might have changed as a result of the treatment (that is, we use the within-person variation we analyzed in Section 4.3). Table 3 presents the most common three-word phrases (trigrams) in these open-ended responses (two-word phrases, or bigrams, turned out not to be very informative).[22] Interestingly, a large share claim that the treatment condition did *not* change their decision, when in fact it appears to have done so, suggesting that for some subjects the effect of reference points is not fully conscious. Some of the more common choices invoke broad notions of fairness and not wanting to disappoint ("unfair chang amount" "didnt disappoint person") as well as noting initial allotments ("chang amount told" "told money initi"). A handful invoked the notion of loss specifically ("person lose money"). None express any concern over how it would affect recipients' impression of the redistributor, as we might have anticipated if our results were driven by subjects being concerned with recipients' forming a negative view of them.

In sum, these experimental results suggest that respect for reference points may help to explain why individuals eschew complete redistribution *even in the absence of moral hazard*, in contradiction to the prescriptions of the optimal tax model under utilitarianism that has been noted by Saez and Stantcheva (2016) and Weinzierl (2014). We show that when redistributors need not consider recipients' reference points, as in our control condition, we essentially recover the Mirrleesian full-redistribution result, making the results from the control arm interesting in their own right. We can thus rehabilitate the full-redistribution result within classic utilitarianism, albeit with *non-standard utility functions.*

[22]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 3-word (trigram) sequences in the remaining text, and calculate frequencies across subject responses.

Table 3: Most common trigrams found in responses when subjects were asked to explain how the treatment condition changed their thinking

didnt chang think	42	give person money	4
didnt chang amount	9	person lose money	4
peopl told money	6	told didnt chang	4
transfer money peopl	6	told receiv amount	4
unfair chang amount	6	werent told made	4
didnt disappoint person	5	chang amount money	3
told amount money	5	chang think felt	3
told money initi	5	chang think made	3
amount money didnt	4	chang think peopl	3
chang amount told	4	chang think person	3
chang think want	4	coin flip decid	3
didnt chang mind	4	didnt money person	3
didnt chang opinion	4	didnt transfer money	3
didnt want fair	4	knew didnt chang	3
felt unfair money	4	make money equal	3

Notes: The exact question wording is: "For the questions where you had to transfer money between other people, how did it change your thinking if these people had already been told how much money they were initially given?" This tables uses data from all MTurk sessions (the higher cost of the *Understanding America Study* forced us to limit the length of the survey, so we were not able to include this question). We include only those observations who gave *different* answers when confronted with the scenario where Persons A and B were not told their initial endowment amounts versus when they were. 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 3-word (trigram) sequences in the remaining text, and calculate frequencies across subject responses.

5 Income-tax survey experiment

Our findings thus far suggest that subjects are sensitive to others' reference points in redistributive decisions in laboratory settings over small stakes. To relate our findings more directly to policy-relevant questions, we now turn to results from a survey experiment on preferences over income tax rates.

The appropriate income tax rate on well-off households is a much-discussed issue in American politics. A threshold of $250,000 has become a focal point, and surveys often ask about support for higher taxes on households with annual incomes of at least that level.[23] We similarly use this threshold in our experiment's design, described below.

5.1 The survey experiment

Subjects were presented with a vignette describing an individual who had received an unexpected increase in earnings. In most waves, the source of the increase was a corporate takeover of the individual's employer (the "takeover" vignette). Subjects were randomly assigned to either a treatment or control arm, which differed only in the *timing* of when the earnings increase took place.[24]

The "control" vignette took the following form:[25]

> There has been much talk about whether wealthy families are paying their fair share in taxes.
>
> Consider the following person. He has been working for about five years as a regional sales manager at a medium-sized firm. *This year*, his firm was taken over by a larger corporation. While he will be doing the same job as before, to

[23] See, e.g., `http://politicalticker.blogs.cnn.com/2012/12/06/trio-of-polls-support-for-raising-taxes-on-wealthy/`.

[24] A version of our survey in which the tax experiment is presented first can be taken at: `https://az1.qualtrics.com/SE/?SID=SV_0MnchCiPWRxAsqV&Preview=Survey&BrandID=columbia`. We do not collect any identifying data—such as IP addresses—via this link.

[25] In the vignettes we reference a tax rate of 22 percent on the "average American," based on NBER Taxsim estimates for combined federal and state income tax, and then add the employee side of payroll taxes.

> make his pay compatible with the earnings of employees in his position at the larger firm, his salary is now doubled, to $250,000.
>
> If it were up to you, how much of his salary should he pay in taxes? (As a basis of comparison, the average American pays about 22 percent in taxes on the income they make.)

In the treatment variant, we attempt to make the protagonists' reference income of $250,000 more deeply embedded. Instead of receiving the raise just this year, he received it five years ago. Specifically, the second paragraph in the treatment vignette reads:

> Consider the following person. He started five years ago as a regional sales manager at a medium-sized firm. *Soon after starting*, his firm was taken over by a larger corporation. While he did the same job as before, his salary was doubled to make his pay compatible with the earnings of employees in his position at the larger firm. Since then, his annual salary has been roughly steady and is now $250,000.

After reading either the control or treatment version of the vignette, subjects provided their response using a slider positioned immediately below the vignette, with values in the range [0,100] percent and the default set to zero. See Appendix Figures A.4 and A.5 for screen shots.

In a later session, we changed the reason for the individual's increase in income. In the control version, the second paragraph of the vignette above is replaced with:

> Consider the following person. *This year*, he won the state lottery. As a result, he will receive $250,000 a year for the rest of his life (note that lottery winnings are treated as taxable income).

As before, in the treatment version we simply replaced "*This year*" with "*Five years ago*" and changed the verb tense as appropriate.

Our implicit assumption is that the five-year reference point in the treatment condition will be viewed by subjects as more strongly embedded and worthy of consideration than the one-year reference point in the control condition. Thus, unlike the main experiment, in which the control group is uninformed of their initial endowments and thus has no ability to form *any* reference point, in this vignette experiment we test whether reference points that are plausibly stronger are respected more by subjects

5.2 Results

We present results in Appendix Table A.6. Across all specifications, subjects choose a lower average tax rate for the protagonist in the five-year vignette, though in the lottery vignette the effect is smaller and statistically insignificant. In the pooled sample, the treatment effect is between 1.2 and 1.7 percentage points, which is a meaningful difference relative to, say, the difference between 2012 Romney and Obama supporters (the former group chooses an average tax rate 2.96 percentage points below the latter).

We relegate these results to the Appendix because, while consistent with our reference-point mechanism, we cannot fully discount other interpretations. Some plausible confounders push *against* finding our result. For example, in the five-year (treatment) scenario, the protagonist would have had more time to accumulate *wealth* and thus could cover the costs of a greater tax burden more easily than the protagonist who only just received the raise (control scenario). Moreover, respondents might think it unfair that, purely due to luck, in the five-year (treatment) scenario the protagonist enjoys the large raise after having barely worked for the company, whereas in the one-year (control) scenario he put in his time before getting the raise. Given the greater willingness to redistribute gains due to luck both in our $5/$15 experiment as well as in work cited earlier, these factors should lead respondents to choose a *higher* tax rate in the control scenario.

However, other potential confounds push in the opposite direction and thus could serve as alternative explanations to our preferred reference-point mechanism. First, subjects confronted with the five-year (treatment) scenario may credit the vignette's protagonist with greater merit because he has worked at the larger (higher paying) corporation for longer. Second, the subject may take into account potential consumption commitments of the protagonist, instead of reference points *per se*. While we could eliminate this concern in the controlled setting of the lab experiment in the previous section, subjects responding to the real-world tax vignette might plausible assume that in the five-year scenario the protagonist chose to purchase a large house or send his children to private school as a result of the higher salary, investments that would be financially or psychologically costly to unwind.

Given these concerns, we view the tax-vignette-experiment result as a first step in testing the real-world implications of our reference-point experiment. We offer additional ideas for future work in the Conclusion.

6 Discussion

It is beyond the scope of the current paper to delve fully into the normative implications of voters or social planners accommodating others' reference points. In this section, we highlight a number of issues that arise from our findings, in the spirit of motivating further work on the subject.

While a voluminous body of work in economics and psychology documents the predictive power of reference-dependence on individual behavior, "the vast literature on reference dependence has virtually entirely avoided welfare analysis" (Reck and Seibold, 2021). In assessing explanations for this gap in the literature, O'Donoghue and Sprenger (2018) write: "[p]erhaps first and foremost is the question of whether gain-loss utility should be given normative weight—i.e., whether we should assume that the same preferences that rationalize behavior should also be used for welfare analysis." In other words, while empirical evidence shows that *individuals* maximize reference-dependent preferences when making decisions, should the *social planner* use these same preferences as an input into the social-welfare-maximization problem? If the social planner gives reference-dependent preferences normative weight, then the answer is yes. But if the social planner believes that reference-dependent preferences represent mistakes or naiveté on the part of the individual (or as Reck and Seibold, 2021 put it: that these preferences "distort behavior relative to what is welfare-maximizing"), then she should not use those preferences as inputs into social-welfare maximization. Instead, paternalistically, the social-planner would maximize welfare under the "correct" preferences subject to the constraint that individual behavior is governed under "incorrect" reference-point-driven preferences.[26]

We believe the results in our paper further complicate the welfare implications of reference dependence because we explore a setting involving inequality and redistribution, and thus implicitly move from a single-agent to a multi-agent setting. In our simple, two-person "society," if the social planner gives normative weight to the reference-point of the richer person, then that decision hurts an already worse-off third party (the poorer person) relative to giving the rich person's reference point no weight.

To underscore the moral ambiguity of applying reference-point arguments to welfare analysis, we note they have been used to argue against the righting of grave historical wrongs. Gunnar Myrdal, in his study of U.S. race relations in the 1940s, found that white Southerners used reference-point-based reasoning to justify why they need not redistribute to (much poorer) Blacks. "[I]t is said that the Negro is accustomed to live on little. . . .'Negroes don't have the same demands on life as white people.' 'They are satisfied with less.' [P]eople accustomed to suffer from want do not feel poverty so much as if they had seen better days....[P]eople who have seen better days are believed to be worse off than other paupers." (Myrdal, 1944) Interestingly, he notes that this reference-point justification is given *in private*, even among educated white Southerners, but not in print, as perhaps whites understood it had more of an intuitive than a logical appeal. Myrdal writes that while "it now seldom gets into respectable print," justifying Jim Crow institutions with the alleged lower reference points of blacks "is widespread in the South and constitutes a most important rationalization among even educated people."[27]

---

[26]In related work extending social-welfare analysis to settings with "behavioral" agents, Farhi and Gabaix (2015) and Lockwood (2015) have calibrated how optimal income tax results may change when agents are susceptible to inattention, present-bias or mental accounting. There is also a somewhat older literature that examines optimal tax results when agents care about their income *relative to others* (see Boskin and Sheshinski, 1978 and Oswald, 1983).

[27]All of the passages we quote from Myrdal are taken from Chapter 4.

A similar set of concerns arise in considering intergenerational relative mobility, which political philosophers tend to view positively as a measure of a society's equality of opportunity (Roemer and Trannoy, 2015). But assuming that adults' reference consumption level is formed in part by the consumption they enjoyed as children via their parents' income, then the attractiveness of an intergenerationally mobile society is called into question. The disutility of the poor adult who grew up rich could outweigh the utility of the rich adult who grew up poor.

Will incorporating reference points into the welfare analysis of redistribution always justify preserving the inequities of the status quo? While beyond the scope of our project, we suspect that this interpretation depends crucially on a static, one-period setting. Below we sketch a few directions future work might take in analyzing the welfare implications of redistribution with reference-point utility in a multi-period setting.

The first way that dynamics can enter into the problem is that the initial disutility of the rich person falling below her reference point should fade with time. That is, if redistribution takes place in period t, then in period $t+1$ her reference point may plausibly be modeled as a blended average of the higher period-t consumption and the lower period-$t+1$ consumption (see, e.g., Di Tella *et al.*, 2010 and DellaVigna *et al.*, 2017, who both show empirical evidence of reference-point adjustment over time). Or if the social-planner can credibly announce that higher levels of redistribution are permanent, then a rich person should adjust downward her expectations of future consumption, again reducing the disutility of future redistribution (see Kőszegi and Rabin, 2006 on expectations-based reference points). With enough time periods and a sufficiently small discount rate, a social planner who gave reference points full normative weight could still engage in substantial redistribution. (The argument in favor of redistribution is further bolstered if individuals underestimate their ability to adjust to new consumption levels, as suggested by the literature on affective forecasting; see, in particular, Kermer *et al.*, 2006 who go so far as to label loss aversion as largely resulting from affective forecasting errors.)

Second, a dynamic model could allow reference points to gain salience as they persist over time (as suggested by DellaVigna *et al.*, 2017, who present evidence for backward-looking reference point determination). A social-planner in a multi-period problem would then have a strong incentive to redistribute sooner rather than later. The longer she waits, the more the rich become accustomed to a high level of post-tax income or consumption and the harder it will be for her to redistribute in the future without lowering the utility of the rich via stronger reference-point attachment. Thus, consideration of the reference points of the rich might push a social planner to aggressively reduce inequality via redistribution as soon as it arises.

7 Conclusion

We provide experimental evidence that individuals who are given the opportunity to redistribute between two recipients with unequal endowments are highly sensitive to the recipients' reference points. When the recipients do *not* know their initial endowments, the redistributor erases nearly the full ex-ante income gap. Recipients knowing their initial endowments reduces redistribution by nearly twenty percent. This effect size is large, nearly half of the effect of having endowments determined via merit rather than luck.

These findings have implications for models of optimal taxation. If losses—even for the wealthy—loom larger than gains, part of the welfare gain from redistribution may be erased. If individuals project their own loss aversion onto others when forming their redistributive preferences, then loss aversion might help explain the gap between voters' stated policy preferences and the more egalitarian normative prescriptions of optimal tax models or the positive predictions from standard political economy models.

People can form reference points over many aspects of their lifestyle, and future work may wish to examine whether third parties respect reference points in other domains. While we have focused on redistribution and the reference points of richer individuals and households, transfer recipients may form reference points over benefits levels. Indeed, the policy of "grandfathering in" benefits cuts suggests that policy-makers respect these potential reference points. In the criminal justice realm, a related question is whether juries are hesitant to sentence white-collar or other well-off defendants from comfortable backgrounds to prison time because they implicitly internalize the adjustment cost, whereas a defendant with past time in prison would not receive such deference. Of course, as with our tax experiment, great care would be needed to distinguish respect for reference points in these scenarios from competing explanations.

References

Abdellaoui, M., Bleichrodt, H. and Paraschiv, C. (2007). Loss aversion under prospect theory: A parameter-free measurement. *Management Science*, **53** (10), 1659–1674.

Alesina, A. and Angeletos, G.-M. (2005). Fairness and redistribution. *American Economic Review*, pp. 960–980.

Allcott, H., Mullainathan, S. and Taubinsky, D. (2014). Energy policy with externalities and internalities. *Journal of Public Economics*.

Almås, I., Cappelen, A. W. and Tungodden, B. (2020). Cutthroat capitalism versus cuddly socialism: Are americans more meritocratic and efficiency-seeking than scandinavians? *Journal of Political Economy*, **128** (5), 1753–1788.

Andreoni, J. and Bernheim, B. D. (2009). Social image and the 50 50 norm: A theoretical and experimental analysis of audience effects. *Econometrica*, **77** (5), 1607–1636.

Ariely, D., Bracha, A. and Meier, S. (2009). Doing good or doing well? image motivation and monetary incentives in behaving prosocially. *American Economic Review*, **99** (1), 544–55.

— and Norton, M. I. (2011). Building a better america—one wealth quintile at a time. *Perspectives on Psychological Science*, **6** (1), 9–12.

Barber IV, B. S. and English, W. (2019). The origin of wealth matters: Equity norms trump equality norms in the ultimatum game with earned endowments. *Journal of Economic Behavior & Organization*, **158** (C), 33–43.

Benabou, R. and Ok, E. (2001). Social mobility and the demand for redistribution: The POUM hypothesis. *Quarterly Journal of Economics*, **116** (2), 447–487.

Bernheim, B. D. and Rangel, A. (2004). Addiction and cue-triggered decision processes. *The American Economic Review*, **94** (5), 1558–1590.

Boskin, M. J. and Sheshinski, E. (1978). Optimal redistributive taxation when individual welfare depends upon relative income. *The Quarterly Journal of Economics*, pp. 589–601.

Chapman, J., Snowberg, E., Wang, S. and Camerer, C. (2018). *Loss attitudes in the US population: Evidence from dynamically optimized sequential experimentation (DOSE)*. Tech. rep., National Bureau of Economic Research.

Charité, J., Fisman, R. and Kuziemko, I. (2015). *Reference Points and Redistributive Preferences: Experimental Evidence*. Tech. rep., National Bureau of Economic Research.

Chetty, R. and Szeidl, A. (2007). Consumption commitments and risk preferences. *The Quarterly Journal of Economics*, **122** (2), 831–877.

Dai, X. (2011). Optimal taxation under income uncertainty. *Annals of Economics and Finance*, **12** (1), 121–138.

Dana, J., Weber, R. A. and Kuang, J. X. (2007). Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness. *Economic Theory*, **33** (1), 67–80.

DellaVigna, S., Lindner, A., Reizer, B. and Schmieder, J. F. (2017). Reference-dependent job search: Evidence from hungary. *The Quarterly Journal of Economics*, **132** (4), 1969–2018.

Dennis, S., Goodson, B. and Pearson, C. (2018). Virtual private servers and the limitations of ip-based screening procedures: Lessons from the mturk quality crisis of 2018.

Di Tella, R., Haisken-De New, J. and MacCulloch, R. (2010). Happiness adaptation to income and to status in an individual panel. *Journal of Economic Behavior & Organization*, **76** (3), 834–852.

Diamond, P. and Saez, E. (2011). The case for a progressive tax: from basic research to policy recommendations. *The Journal of Economic Perspectives*, **25** (4), 165–190.

Dimick, M., Rueda, D. and Stegmueller, D. (2016). The altruistic rich? inequality and other-regarding preferences for redistribution. *Quarterly Journal of Political Science*, **11** (4), 385–439.

—, — and — (2018). Models of other-regarding preferences, inequality, and redistribution. *Annual Review of Political Science*, **21**, 441–460.

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

—, — and Van der Weele, J. J. (2013). Preferences for redistribution and perception of fairness: An experimental study. *Forthcoming, Journal of the European Economic Association.*

Engström, P., Nordblom, K., Ohlsson, H. and Persson, A. (2015). Tax compliance and loss aversion. *American Economic Journal: Economic Policy*, **7** (4), 132–64.

FARHI, E. and GABAIX, X. (2015). *Optimal taxation with behavioral agents*. Tech. rep., National Bureau of Economic Research.

FISMAN, R., GLADSTONE, K., KUZIEMKO, I. and NAIDU, S. (2020). Do americans want to tax wealth? evidence from online surveys. *Journal of Public Economics*, **188**, 104207.

—, JAKIELA, P. and KARIV, S. (2017). Distributional preferences and political behavior. *Journal of Public Economics*, **155**, 1–10.

GÄCHTER, S. and RIEDL, A. (2005). Moral property rights in bargaining with infeasible claims. *Management Science*, **51** (2), 249–263.

HALLSWORTH, M., LIST, J. A., METCALFE, R. D. and VLAEV, I. (2017). The behavioralist as tax collector: Using natural field experiments to enhance tax compliance. *Journal of Public Economics*, **148**, 14–31.

JONES, P. (2020). Loss aversion and property tax avoidance. *Available at SSRN 3511751*.

KAHNEMAN, D. and TVERSKY, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the Econometric Society*, pp. 263–291.

KANBUR, R., PIRTTILÄ, J. and TUOMALA, M. (2008). Moral hazard, income taxation and prospect theory. *Scandinavian Journal of Economics*, **110** (2), 321–337.

KENNEDY, R., CLIFFORD, S., BURLEIGH, T., WAGGONER, P. and JEWELL, R. (2018). The shape of and solutions to the mturk quality crisis. *Available at SSRN*.

KERMER, D. A., DRIVER-LINN, E., WILSON, T. D. and GILBERT, D. T. (2006). Loss aversion is an affective forecasting error. *Psychological science*, **17** (8), 649–653.

KUZIEMKO, I., NORTON, M. I., SAEZ, E. and STANTCHEVA, S. (2015). How elastic are preferences for redistribution? evidence from randomized survey experiments. *American Economic Review*, **105** (4), 1478–1508.

KŐSZEGI, B. and RABIN, M. (2006). A Model of Reference-Dependent Preferences*. *The Quarterly Journal of Economics*, **121** (4), 1133–1165.

LEE, W. and ROEMER, J. E. (2006). Racism and redistribution in the United States: A solution to the problem of American exceptionalism. *Journal of Public Economics*, **90** (6), 1027–1052.

LEVENTHAL, G. S. (1980). *What should be done with equity theory?* Springer.

LIST, J. (2007). On the interpretation of giving in dictator games. *Journal of Political Economy*, **115**, 482–493.

Lockwood, B. (2015). Optimal taxation with present bias.

Lü, X. and Scheve, K. (2016). Self-centered inequity aversion and the mass politics of taxation. *Comparative Political Studies*, **49** (14), 1965–1997.

Meltzer, A. and Richard, S. (1981). A rational theory of the size of government. *The Journal of Political Economy*, **89** (5), 914–927.

Mirrlees, J. A. (1971). An exploration in the theory of optimum income taxation. *Review of Economic Studies*, **38** (2), 175–208.

Myrdal, G. (1944). *An American Dilemma: The Negro Problem and Modern Democracy*. Harper & Row.

O'Donoghue, T. and Rabin, M. (2006). Optimal sin taxes. *Journal of Public Economics*, **90** (10), 1825–1849.

— and Sprenger, C. (2018). Reference-dependent preferences. In *Handbook of Behavioral Economics: Applications and Foundations 1*, vol. 1, Elsevier, pp. 1–77.

Oswald, A. J. (1983). Altruism, jealousy and the theory of optimal non-linear taxation. *Journal of Public Economics*, **20** (1), 77–87.

Reck, D. and Seibold, A. (2021). The welfare economics of reference dependence.

Rees-Jones, A. (2017). Quantifying loss-averse tax manipulation. *The Review of Economic Studies*, **85** (2), 1251–1278.

Roemer, J. E. and Trannoy, A. (2015). Chapter 4 - equality of opportunity. In A. B. Atkinson and F. Bourguignon (eds.), *Handbook of Income Distribution*, *Handbook of Income Distribution*, vol. 2, Elsevier, pp. 217–300.

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

Trump, K.-S. (2015). Accepting inequality: How "what is" influences "what ought to be".

Weinzierl, M. (2014). The promise of positive optimal taxation: normative diversity and a role for equal sacrifice. *Journal of Public Economics*, **118**, 128–142.

Online Appendix

A Additional figures and tables

Figure A.1: Redistribution with classical and reference-dependent utility functions

(a) Utility for “poor” individual p, classical utility

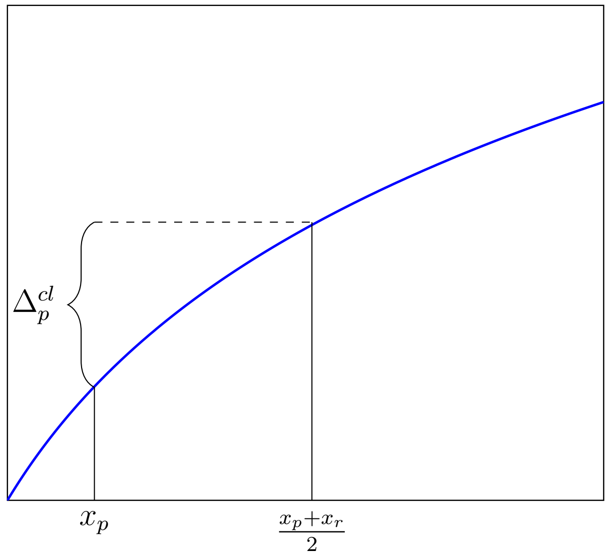

(b) Utility for “rich” individual r, classical utility

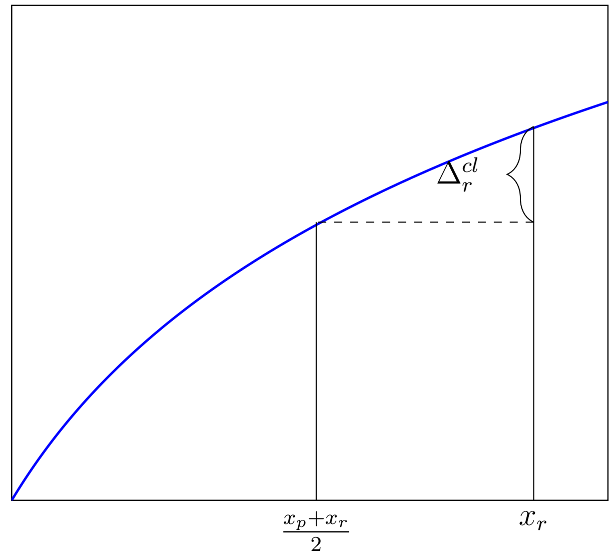

(c) Utility for “poor” individual p, reference-dependent utility

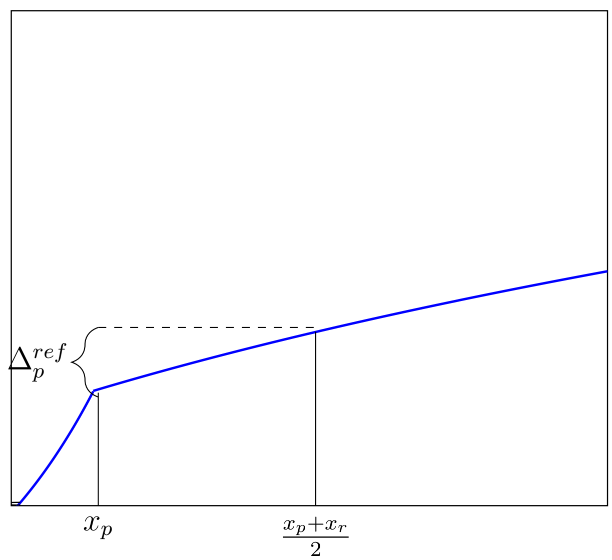

(d) Utility for “rich” individual r, reference-dependent utility

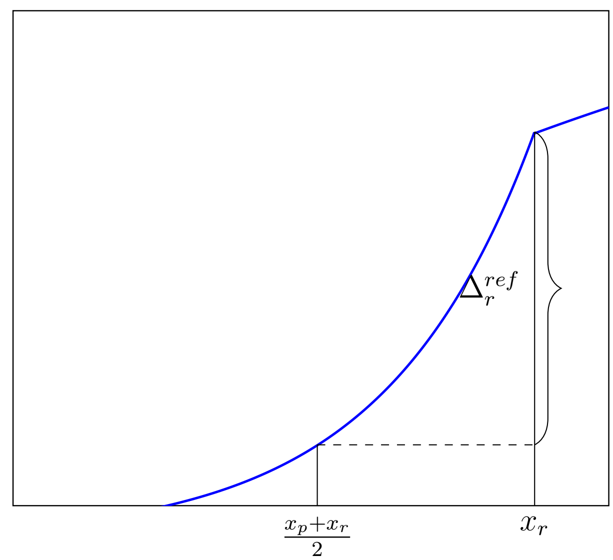

Notes: The figures are drawn by the authors for the sake of illustration. The two top panels show the classic optimal-tax result when initial endowments are exogenous and utility is identical and with standard properties. Under these assumptions, the loss to the rich is smaller than the gain to the poor ($-\Delta_r^{cl} < \Delta_p^{cl}$) and in fact the welfare-maximizing allocation is full equalization at $\frac{x_p+x_r}{2}$. The bottom two panels show how this result can break down under reference-dependent utility, as full redistribution yields lower total utility than the status-quo unequal endowments ($-\Delta_r^{ref} > \Delta_p^{ref}$).

Figure A.2: Screenshot of the main redistribution experiment (control arm)

Consider two other participants on mTurk, person A and person B. Based on a coin flip, we have given $5 to person A and $15 to person B.

You can now transfer money between persons A and B. *Persons A and B are not told how much money they were initially given.* If you decide to give Person A $X instead of $5, they will simply be told that they have been given $X, and will not know how much they started with. *Nor* will they know that there is another person (Person B) involved, or that a third party (you) determined the money they received.

Please indicate below what transfer, if any, you would make.

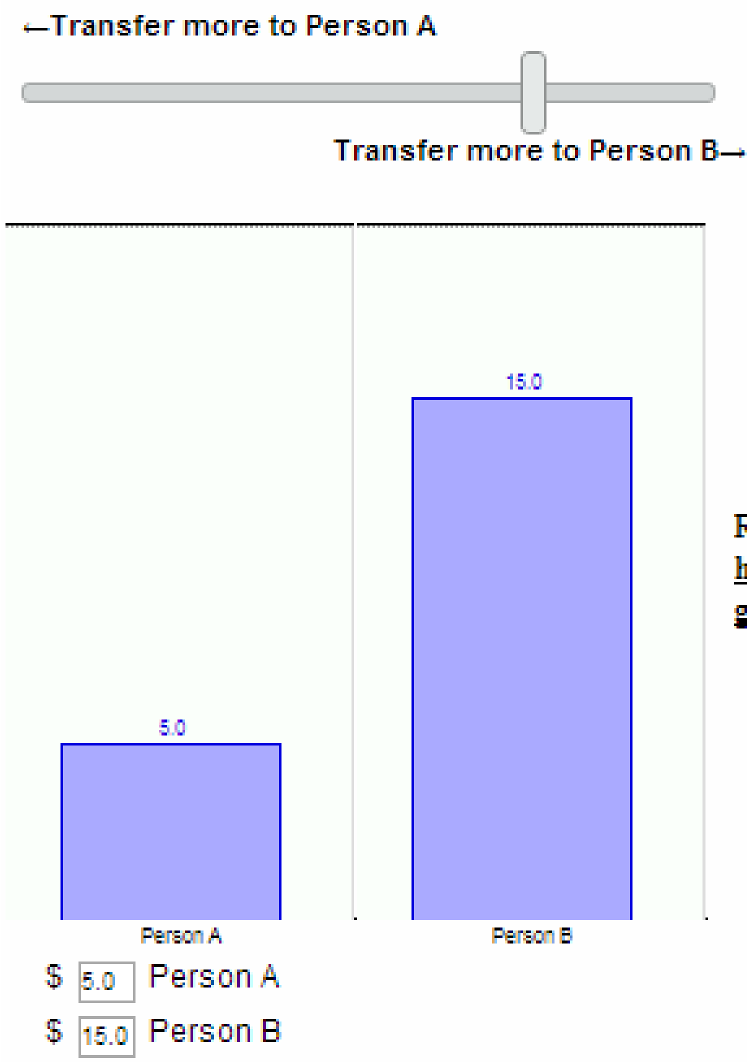

Recall: Person A and B do NOT know how much money they were initially given.

Figure A.3: Screenshot of the main redistribution experiment (treatment arm)

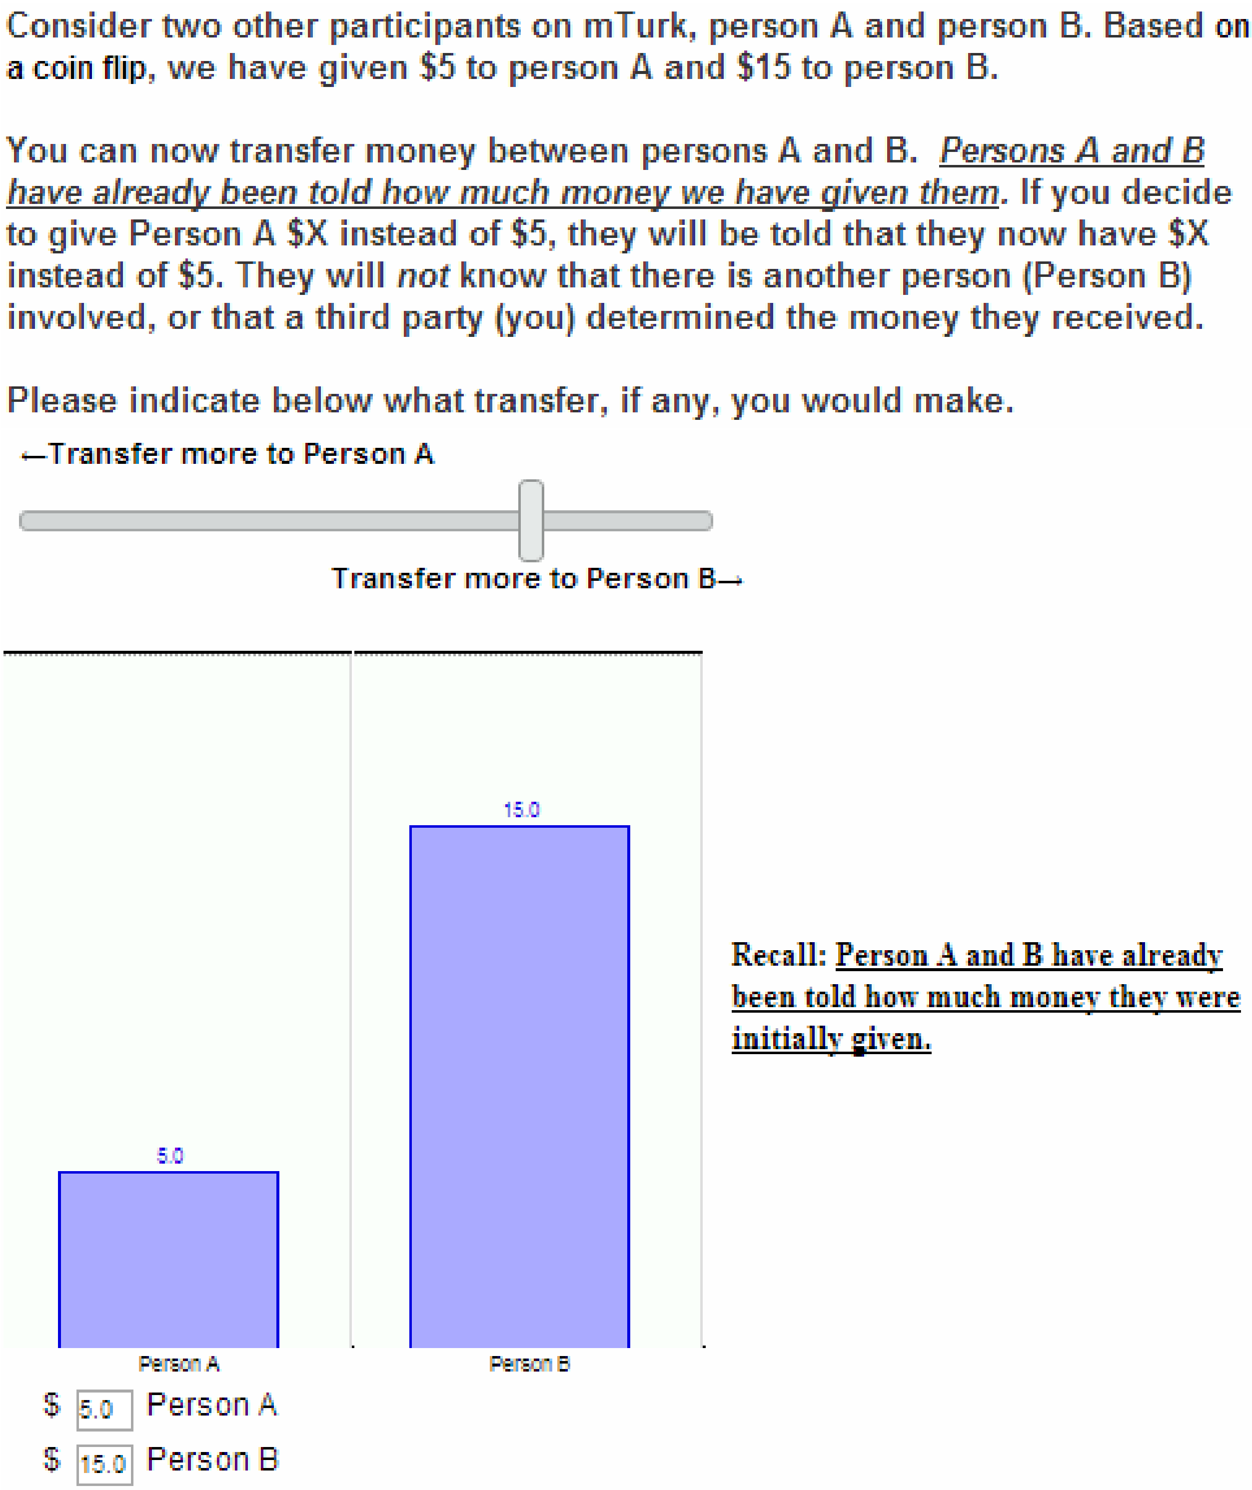

Figure A.4: Screenshot of income-tax experiment (control arm)

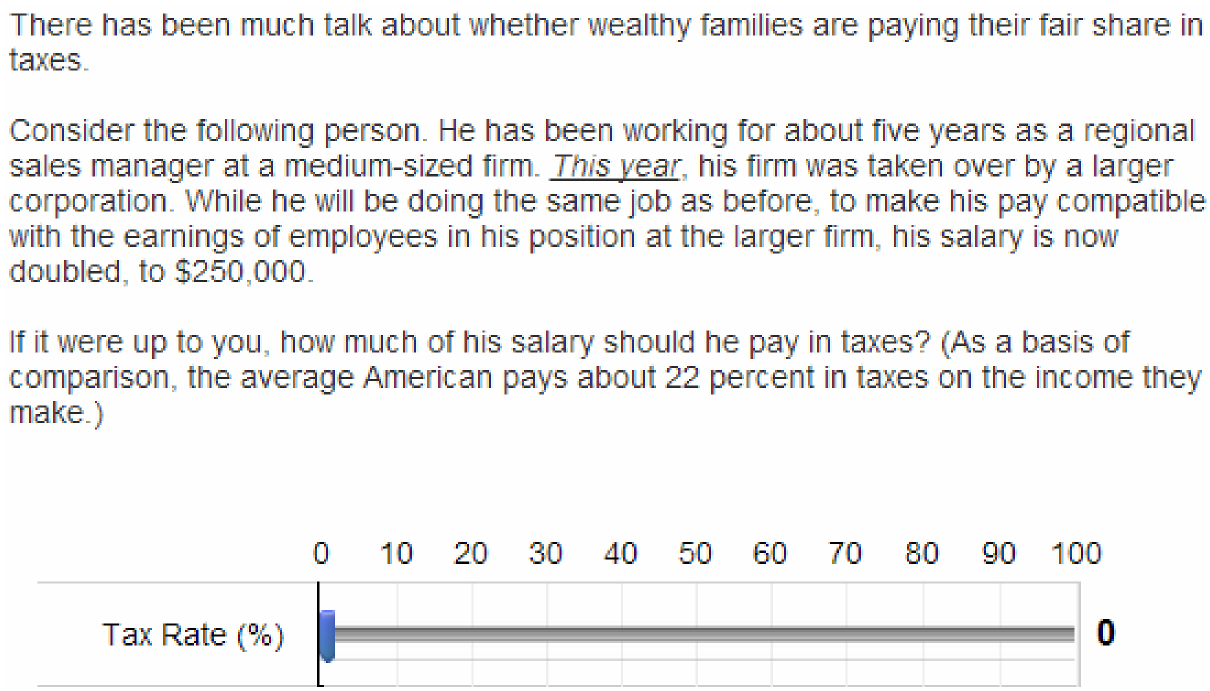
There has been much talk about whether wealthy families are paying their fair share in taxes.

Consider the following person. He has been working for about five years as a regional sales manager at a medium-sized firm. *This year*, his firm was taken over by a larger corporation. While he will be doing the same job as before, to make his pay compatible with the earnings of employees in his position at the larger firm, his salary is now doubled, to $250,000.

If it were up to you, how much of his salary should he pay in taxes? (As a basis of comparison, the average American pays about 22 percent in taxes on the income they make.)

Figure A.5: Screenshot of income-tax experiment (treatment arm)

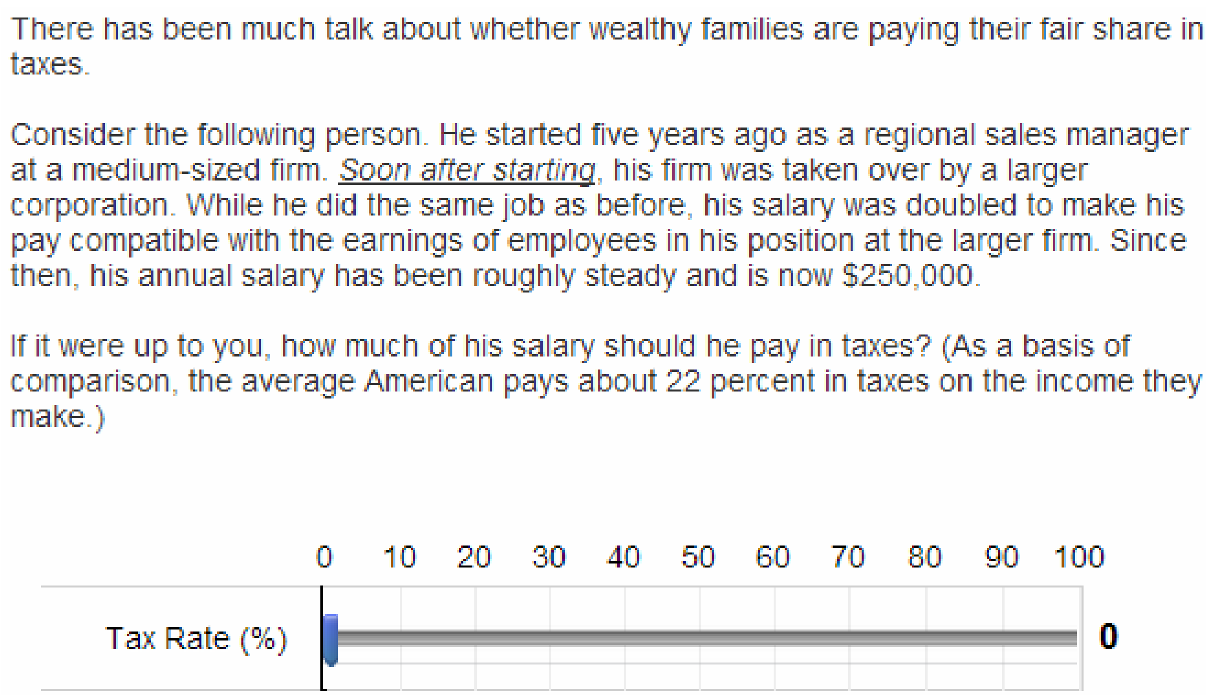
There has been much talk about whether wealthy families are paying their fair share in taxes.

Consider the following person. He started five years ago as a regional sales manager at a medium-sized firm. *Soon after starting*, his firm was taken over by a larger corporation. While he did the same job as before, his salary was doubled to make his pay compatible with the earnings of employees in his position at the larger firm. Since then, his annual salary has been roughly steady and is now $250,000.

If it were up to you, how much of his salary should he pay in taxes? (As a basis of comparison, the average American pays about 22 percent in taxes on the income they make.)

Table A.1: Survey Session Details

Date	Observations Total	Unique	First Exper.	Version of $5/$15 exp.	Income tax exp.?	Platform
Reference point experiments						
Feb 13, 2014	187	187	$5/$15	Standard	No	MTurk
Feb 27, 2014	312	295	$5/$15	Standard	No	MTurk
Mar 10, 2014	301	250	$5/$15	Standard	No	MTurk
Mar 21, 2014	352	282	Tax	Standard	Yes	MTurk
Mar 24, 2014	374	303	Tax	Standard	Yes	MTurk
May 28, 2014	312	207	$5/$15	Real Money	Yes	MTurk
May 30, 2014	332	216	$5/$15	No Emphasis	Yes	MTurk
Jun 19, 2014	314	200	Tax	Standard	Yes	MTurk
Dec 11, 2014	307	196	$5/$15	No Promises	No	MTurk
June–July 2019	308	308	n.a.	Standard	No	UAS
Luck v. merit experiment						
Apr 25, 2014	321	228	$5/$15	Luck/Merit	No	MTurk

Notes: Total observations and analysis sample observations differ because in almost all analyses we drop anyone who took a previous survey. We collect data on two different platforms: MTurk (Amazon Mechanical Turk) and the UAS (Understanding America Study) panel. On MTurk, each session ran for one or two hours. On UAS, the data were collected from June 24 to July 19, 2019. Further details on the wording used in each session can be found in the text (Sections 2 and 4).

Table A.2: Assessing perceptions of bias and survey fatigue, by redistribution experiment survey arm

	LW bias	RW bias	No bias	Minutes
	(1)	(2)	(3)	(4)
Treated in first stage	0.0094	-0.0003	-0.0132	0.1479
	(0.0143)	(0.0093)	(0.0168)	(0.3070)
Cont. gp. mean	0.103	0.042	0.850	12.257
Observations	1883	1883	1883	1733

Notes: The outcome variables in the first two columns is a binary variable for the subject reporting they felt that the survey experiment exhibited left-wing or right-wing bias, respectively. Subjects who finished the survey very quickly were not included in these regressions. Also, we did not ask the bias questions to the UAS sample and thus they are also excluded.

Table A.3: Basic summary statistics in MTurk and UAS sample compared to the General Social Survey

	MTurk	UAS	GSS
Age	33.03	53.23	47.44
Female	0.44	0.57	0.55
White	0.77	0.87	0.73
Black	0.07	0.08	0.15
Hispanic	0.06	0.08	0.16
Asian	0.08	0.03	
Has at least college education	0.45	0.45	0.32
Subject Income (Thousands of Dollars)	49.53	72.96	71.56
Supported Obama in 2012	0.64	.	0.62
Supported Clinton in 2016	.	0.48	0.55
Observations	2029	315	7753

Notes: Col. 1 includes all ten sessions of the reference-point experiment conducted in MTurk. Col. 2 includes the session conducted on the UAS sample. Col. 3 includes all adults in the 2014, 2016, and 2018 GSS (weighted with the provided individual-level weights), except that only the 2014 and 2016 surveys are used for the Obama question and only the 2018 survey is used for the Clinton question. "Income" refers to household income (in units of $1,000)

Table A.4: Further summary statistics and experimental balance

	Means				Observations	
	Treat	Control	Diff.	p-val	Treat	Control
Age	35.964	35.712	0.252	0.659	1189	1127
Female	0.460	0.457	0.003	0.882	1189	1127
White	0.792	0.783	0.010	0.570	1189	1127
Black	0.071	0.078	-0.007	0.495	1189	1127
Hisp	0.056	0.060	-0.004	0.682	1189	1127
Asian	0.073	0.067	0.007	0.532	1189	1127
Income (in thousands of dollars)	52.745	52.007	0.738	0.669	1189	1126
Fulltime	0.442	0.430	0.012	0.559	1189	1127
Partime	0.143	0.123	0.020	0.165	1189	1127
College	0.451	0.445	0.005	0.795	1189	1127
Supported Dem. in last pres. election	0.628	0.620	0.007	0.724	1144	1072

Notes: The first column displays means for those randomized into the treatment version of the $5/$15 money-transfer experiment (where recipients do know their original endowment) and the second column displays the means for the control version (where recipients do not know their original endowment). The third column displays the difference and the fourth column the p-value associated with $H_0 : Diff = 0$ for the difference between treatment and control groups. The final two columns give sample sizes for the treatment and control groups. This table includes only the ten sessions where the experiment includes the $5/$15 reference-point experiment. It does not include the $5/$15 luck-merit experiment (analyzed in Table 2). The experimental balance table for this experiment is available upon request. The number of observations differ across variables because of subject non-response.

Table A.5: Within-subject results

	Dept. var: Amount redistributed		
	(1)	(2)	(3)
Treated in first stage	-0.661***	-0.600***	-0.715***
	(0.077)	(0.119)	(0.100)
Dept. var. mean	3.965	3.909	4.016
Sample	All	$T \rightarrow C$	$C \rightarrow T$
Observations	4306	2044	2262

Notes: All regressions include respondent fixed effects so as to make use of within-subject variation (and thus the number of observations are twice the number in the baseline, between-subject specifications). Subjects who finished the survey very quickly were excluded. $C \rightarrow T$ denotes the subsample that was first randomized into the control scenario and *then* the treatment scenario. $T \rightarrow C$ denotes the subsample that was first randomized into the treatment scenario and *then* the control scenario. Standard errors are clustered at the subject level. $^{*}p < 0.1, ^{**}p < 0.05, ^{***}p < 0.01$.

Table A.6: Preferred average tax rate for person who became rich five versus one year ago

	Dept. var: Chosen tax rate						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Treated (rich for five yrs.)	-0.0117*	-0.0171**	-0.0160**	-0.0197**	-0.0156*	-0.0099	-0.0189
	(0.0064)	(0.0075)	(0.0073)	(0.0087)	(0.0082)	(0.0150)	(0.0160)
Cont. gp. mean	0.2877	0.2865	0.2841	0.2904	0.2872	0.2764	0.2764
Ex. presented second	No	Yes	Yes	Yes	Yes	Yes	Yes
Controls	No	No	Yes	No	Yes	No	Yes
Vignette	Both	Both	Both	Takeover	Takeover	Lottery	Lottery
Observations	1097	721	708	532	523	189	185

Notes: The outcome variable is the tax rate chosen for the protagonist in the vignette. "Treated" refers the subject being told that the protagonist's income had increased five years ago (as opposed to the control group, who were told that the protagonist's income increased only one year ago). All regressions include session fixed effects. Subjects who finished the survey very quickly are excluded from the regressions. "Vignette" refers to the brief description of the event that led to the sudden increase in earnings. "Controls" indicates the inclusion of the following covariates: age, female, white, black, Hispanic, Asian, income, student status, full-time status, part-time status, Democratic-candidate support in the most recent presidential election, and college degree. $^{*}p < 0.1, ^{**} p < 0.05, ^{***} p < 0.01$.

Table A.7: Differential treatment effects between whose who took the \$5/\$15 experiment first and those who took it second

	Amount redistributed		Zero Redistribution	
	(1)	(2)	(3)	(4)
Treated	-0.738***	-0.761***	0.104***	0.092***
	(0.155)	(0.180)	(0.024)	(0.027)
Treated × Presented first	-0.092	-0.063	0.055*	0.065*
	(0.207)	(0.223)	(0.032)	(0.033)
Cont. gp. mean	4.374	4.579	0.118	0.081
Ex. short duration	Yes	Yes	Yes	Yes
Controls	No	Yes	No	Yes
Observations	2194	1888	2194	1888

Notes: The outcome in the first two columns is the amount of dollars redistributed from the recipient who starts with \$15 to the recipient with \$5, and the (binary) outcome in the final two columns is whether the subject chose zero redistribution. Treated refers to the subject being told that the recipients knew their initial endowments (as opposed to the control group who were told that the recipients would only know their final allocations). Presented first refers to survey sessions where the main \$5/\$15 reference-point experiment was presented first. Controls indicates the inclusion of the following covariates: age, female, white, black, Hispanic, Asian, income, student status, full-time status, part-time status, Democratic-candidate support in the most recent presidential election, and college degree. All regressions include session fixed effects. Ex. short duration: exclude subjects who finish the survey in a suspiciously short amount of time. $^{*}p < 0.1,^{**} p < 0.05,^{***} p < 0.01$.

B Simulation-based inference for between-subject results

In this section, we formally investigate how many sessions out of ten we would expect to fail to reach statistical significance if the pooled treatment effect we estimate is "the truth." Note that we perform this exercise only for the results in Figure 1, which use the between-subject experimental variation, because the by-session within-subject results in Figure 3 are nearly all significant at the five-percent level.

B.1 Procedure

We proceed as follows:

1. We obtained the estimated treatment effect ($\hat{\beta}$), a set of estimated session fixed effects ($\hat{\mu}_{session(i)}$), and an estimate of the error term's variance ($\hat{s}$) based on results from the pooled regression that we consider to be our main results in the paper (Table 1, column 2).

2. We then generate 10 datasets with sample sizes that reflect those in our actual study. The datasets were constructed as follows:

   (a) All observations i were randomly assigned to treatment or control, to generate the variable $Treat_i$.

   (b) The dependent variable that captures the amount of redistribution in the simulated data, $Simul\_Redist_i$, is obtained using the estimates derived in the first step above. That is, we generate simulated data using the following:

   $$Simul\_Redist_i = \hat{\beta} Treat_i + \hat{\mu}_{session(i)} + \epsilon_i$$

   with $\epsilon_i \sim N(0, \hat{s})$, where $\hat{s}$ is the estimated variance of the error term in the pooled regression in the first step.

3. We then use this *simulated* data to run, for each simulated session, the following specification, which exactly parallels the construction of the figure that provides session-specific treatment effects (Figure 2 in our original submission). That is, we run the following for each of the 10 sessions:

   $$Simul\_Redist_i = b_0 + b_1 Treat_i + u_i$$

We conduct this process 200 times.

B.2 Results

Appendix Figure B.1(a) shows the distribution of the number of sessions with insignificant estimates at the 5% significance level. In all 200 runs, only 3 partitions have significant estimates (at 5% significance level) in all ten sessions. On average, out of the ten sessions, 3.47 treatment effects were not significant at the 5% level. Thus, even when the treatment effect exists in the true model in the overall population, we would expect several insignificant estimates in subpopulations due to sampling error.

We generated comparable simulated results for the extensive-margin measure of redistribution in Appendix Figure B.1(c). This simulation is a bit more complicated, given the discrete nature of the outcome; we employ a linear probability model, with a cutoff of 0.25 in the probability of zero-redistribution. In our 200 simulated datasets, there are on average 4.2 sessions with insignificant treatment effects at the 5% significance level, a result that is largely invariant to our choice of cutoff for assignment to zero-redistribution.

Appendix Figures B.1(b) and (d) are the ten-percent-significance analogues to Appendix Figures (a) and (c), respectively. As we would expect, fewer of the simulated rounds fail at the ten-percent significance level. Nonetheless, we would expect 2.45 failures when overall redistribution is the outcome and 3.14 when zero-redistribution is the outcome.

Figure B.1: Number of sessions with insignificant estimates

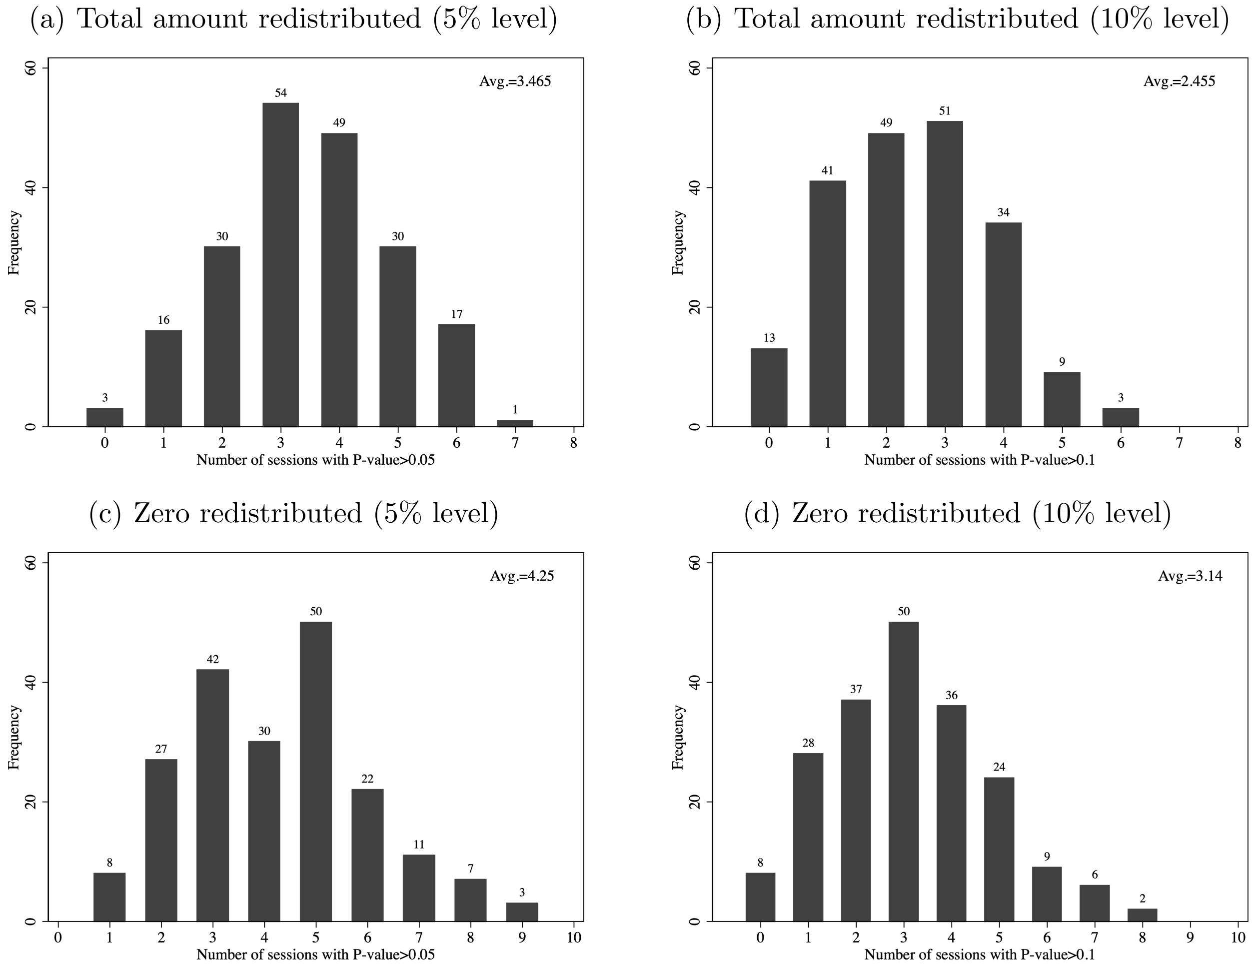

Notes: This figure show the distribution of the number of sessions with p-values greater than 0.05 and 0.1 from 200 simulation partitions. We do the simulations for two outcomes: (a) the amount redistributed and (b) the (binary) outcome of choosing to redistribute \$0.

C Comparison with Almås *et al.* (2020)

C.1 Overview

Almås *et al.* (2020) is probably the closest to our paper in experimental design, in particular a "follow-up" experiment they describe in the paper, the results of which are documented in an online appendix.

The follow-up experiment instructions are as follows, copied from their Appendix. The first is their control ("no info" is provided to the workers about the lottery results, which determine the initial payouts) and the second is their treatment ("info" in the form of the lottery results *are* provided to the workers).

Luck treatment (as in main experiment)

In contrast to traditional survey questions that are about hypothetical situations, we now ask you to make a choice that has consequences for a real life situation. A few days ago two individuals, let us call them worker A and worker B, were recruited via an international online market place to conduct an assignment.

They were each offered a participation compensation of 2 USD regardless of what they were paid for the assignment. After completing the assignment, they were told that their earnings from the assignment would be determined by a lottery. The worker winning the lottery would earn 6 USD for the assignment and the other worker would earn nothing for the assignment. They were not informed about the outcome of the lottery. However, they were told that a third person would be informed about the assignment and the outcome of the lottery, and would be given the opportunity to redistribute the earnings and thus determine how much they were paid for the assignment.

You are the third person and we now want you to choose whether to redistribute the earnings for the assignment between worker A and worker B. Your decision is completely anonymous. The workers will receive the payment that you choose for the assignment within a few days, but will not receive any further information.

Worker A won the lottery and earned 6 USD for the assignment, thus worker B earned nothing for the assignment.

Please state which of the following alternatives you choose:

Luck-info treatment

In contrast to traditional survey questions that are about hypothetical situations, we now ask you to make a choice that has consequences for a real life situation. A few days ago two individuals, let us call them worker A and worker B, were recruited via an international online market place to conduct an assignment.

They were each offered a participation compensation of 2 USD regardless of what they were paid for the assignment. After completing the assignment, they were told that their earnings from the assignment would be determined by a lottery. The worker winning the lottery would earn 6 USD for the assignment and the other worker would earn nothing for the assignment. They were informed about the outcome of the lottery. However, they were also told that a third person would be informed about the assignment and the outcome of the lottery, and would be given the opportunity to redistribute the earnings and thus determine how much they were paid for the assignment.

You are the third person and we now want you to choose whether to redistribute the earnings for the assignment between worker A and worker B. Your decision is completely anonymous. The workers will receive the payment that you choose for the assignment within a few days, but will not receive any further information.

Worker A won the lottery and earned 6 USD for the assignment, thus worker B earned nothing for the assignment.

Please state which of the following alternatives you choose:

The key result is reproduced in Appendix Figure C.1. No effect is observed in the U.S.

sample and a small effect (it appears insignificant) is observed for the Norwegian sample in the same direction as our results, in that there is less redistribution when the "workers" know their original payments than when they do not.

Note that Almas et al. present their results in terms of the Gini coefficient of the final (ex-post) distribution. The Gini for a two-person distribution can be written as:

$$Gini = \frac{|income\ person\ A - income\ person\ B|}{total\ income} \tag{1}$$

where person A is the recipient with higher initial endowments. While in our paper we presented results in terms of dollars redistributed, we adjust our results below to show the ex-post Gini for treatment versus control for ease of comparison.

C.2 Why no reference-point effects in Almås *et al.* (2020)?

There are three key differences in the design of the Almas et al. "follow up" and our reference-point experiment, all of which serve to soften or erase the likely reference points of "workers" in Almas et al. relative to the "recipients" in our experiment.

The first difference is that the payment in Almas et al. *is explicitly described as subject to change*. Quoting from their instructions above: "They [workers] were informed about the outcome of the lottery. However, they were also told that a third person would be informed about the assignment and the outcome of the lottery, and would be given the opportunity to redistribute the earnings and thus determine how much they were paid for the assignment." Thus, in this sense, their set-up is closest to our Procedural Justice variant, in which we *intentionally* try to soften the reference point and indeed find a smaller (thought still marginally significant) treatment effect.

A second important difference is that "workers" in Almas et al. know about each other because it is implicit in the instructions that there is a winner and a loser of the lottery (and that the winner gets $6 and the loser gets $0). So, in their luck-info version, the lottery-loser knows that even though they completed the task, they receive no payment (beyond the show-up fee). And the lottery-winner knows that the loser gets no payment despite having completed the task and that she, the winner, is getting paid instead. Given the unfairness of the situation and that the workers know there is ex-post redistribution by a fully-informed third party, they would likely expect the redistributor to correct at least part of the unfairness. Again, this set-up would serve to weaken the reference point of the higher-paid worker.

By contrast, in our set up (even in our No Promises version), we explicitly tell the redistributor that person A and person B *do not know about each other or about the redistributor*: "They [person A] will not know that there is another person (Person B) involved, or that a third party (you) determined the money they received." Our persons A and B perform no work (they are just given initial, unequal endowments based on a coin toss) and thus there is no inherent unfairness of the unlucky person completing a task without compensation.

The final important difference that serves to weaken the reference point in Almas et al. is that the payment is given "within a few days," which should again soften the reference point. It is very possible the worker would not even remember what payment they had received initially and thus the redistributor could well assume that the higher-paid worker will suffer limited if any sense of loss relative to the reference point.

C.3 Other results in Almås *et al.* (2020)

While less directly relevant to the question of reference points and redistribution, there are other experimental results in Almås *et al.* (2020) that roughly parallel those in our paper.

For example, in both papers, redistribution when luck versus merit drives the initial difference is compared (though merit in our experiment is based on SAT questions and in their experiment it is based on an effort-intensive task). Figure C.2(a) shows the average level of inequality implemented in the treatment (merit) and control (luck) groups. Consistent with our finding, we observe that there is more inequality (less redistribution) in the treatment arm of the experiment. However, the magnitudes of inequality are not as large as those in Almås *et al.* (2020). Specifically, the average implemented inequality in the merit treatment is around 0.27 which is relatively smaller compared to 0.56 in Almås *et al.* (2020). This is not surprising because this measure is quite sensitive to the scale of status quo allocation. The high implemented inequality in their paper can be attributed to the high inequality of the status quo allocation.

To make a comparison between merit effects in our paper and the findings in Almås *et al.* (2020), we calculated the implemented inequality of the luck-vs-merit sample by the same measure used in their paper (as described above) and conduct a parallel exercise. Columns (1) and (2) of Table C.1 report the corresponding regressions of implemented inequality on the treatment. We reproduced the baseline merit effects in our main analysis in Columns (3) and (4). Columns (1) and (2) show that merit is a large deterrent against redistribution: it significantly increases implemented inequality by 0.159 (0.122), which is quite similar to the merit effect (0.195) found in Almås *et al.* (2020).

We also estimate the shares of egalitarians, meritocrats, and libertarians in our luck-vs-merit sample. Following Almås *et al.* (2020), these three fairness views are estimated as follows:

i) Egalitarians: the share of redistributors dividing payoffs equally in the merit treatment.

ii) Meritocrats: the difference between the share of redistributors allocating more to the more productive (well-performed) recipient in the merit treatment and the share distributors allocating more to the lucky recipient in the luck treatment.

iii) Libertarians: the share of redistributors allocating everything to the lucky recipient in the luck treatment.

As shown in Figure C.3, meritocratism is the most prevalent fairness view (about 40.4%) in our luck-vs-merit experiment, which is almost the same as the share in Almås *et al.* (2020) paper (42.5%). Interestingly, we also observe that a larger share of redistributors chose according to an egalitarian fairness view (34%), compared to 15.3% in Almås *et al.* (2020) paper. This may help us explain why the implemented inequality is smaller than what was found in their paper. That is, given the status quo allocation of $5 and $15 dollars, there seem to be more people who consider both luck and merit to be unfair in redistribution. It also raises the possibility that people's fairness views could be affected by the status quo allocation. (Given the nature of the luck-vs-merit experiment, we have shut down any reference-point effect in both the treatment and control arms of the experiment.)

Figure C.1: Comparisions between Almås *et al.* (2020)'s follow-up and our main experiments

(a) Almås *et al.* (2020): United States

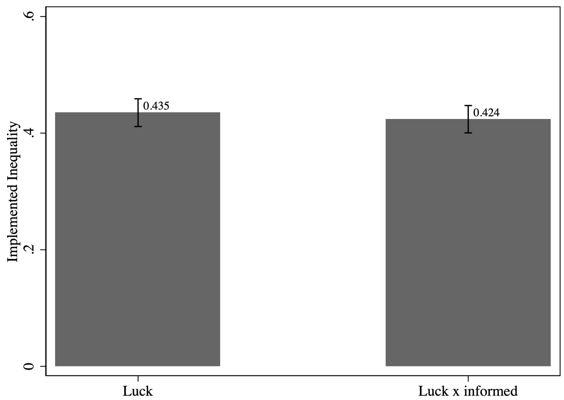

(b) Almås *et al.* (2020): Norway

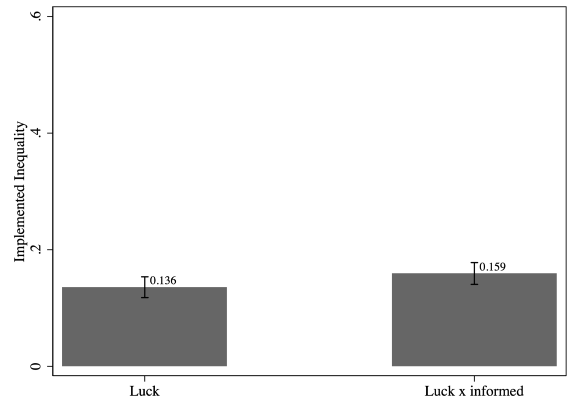

(c) Our paper: Pooled

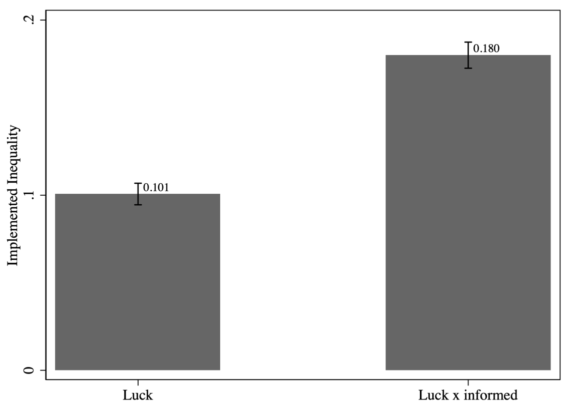

(d) Our paper: Procedural justice

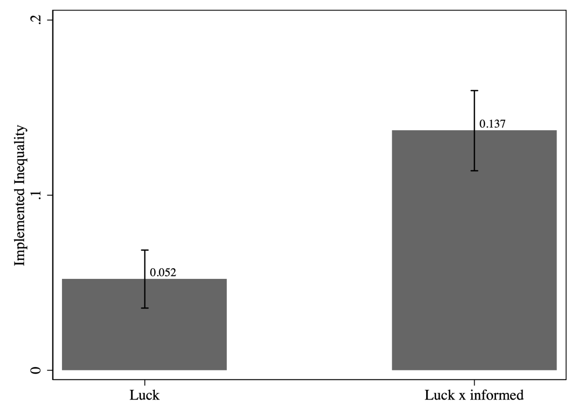

Notes: Figures (a) and (b) replicate the Figure A2 in Almås *et al.* (2020). Figures (c) and (d) show the average level of implemented inequality (as defined in Almås *et al.* (2020)) by pooled and "procedural justice" sessions for "treatment" (Luck x informed) and "control" (Luck) groups.

Table C.1: Using Almas et al. (2020)'s measure to estimate the merit effects

	Implemented inequality		Amount redistributed	
	(1)	(2)	(3)	(4)
Treated (Merit)	0.159***	0.122***	-1.947***	-1.806***
	(0.026)	(0.033)	(0.286)	(0.360)
Control mean	0.116	0.126	4.428	4.515
Ex. short duration	No	Yes	No	Yes
Ex. repeat-takers	No	Yes	No	Yes
Sample	SAT/coin	SAT/coin	SAT/coin	SAT/coin
Observations	311	206	311	206

Notes: The outcome in the first two columns is the implemented inequality calculated by the same measure used in Almas et al. (2020), while the outcome in the last two columns is the number of dollars redistributed from the recipient who starts with \$15 to the recipient with \$5. "Treated" in this luck-vs-merit experiment refers to the subject being told that the recipients were assigned their initial endowments based on performance on SAT-type questions. The control group is told that the recipients do not know their initial endowments and that the endowments were determined by a coin flip. "Ex. short duration": exclude subjects who finish the survey in a suspiciously short amount of time. "Ex. repeat takers": exclude repeat-talers in the sample. $^{*}p < 0.1,^{**}p < 0.05,^{***}p < 0.01$.

Figure C.2: Average implemented inequality in our paper

(a) Average implemented inequality in our paper

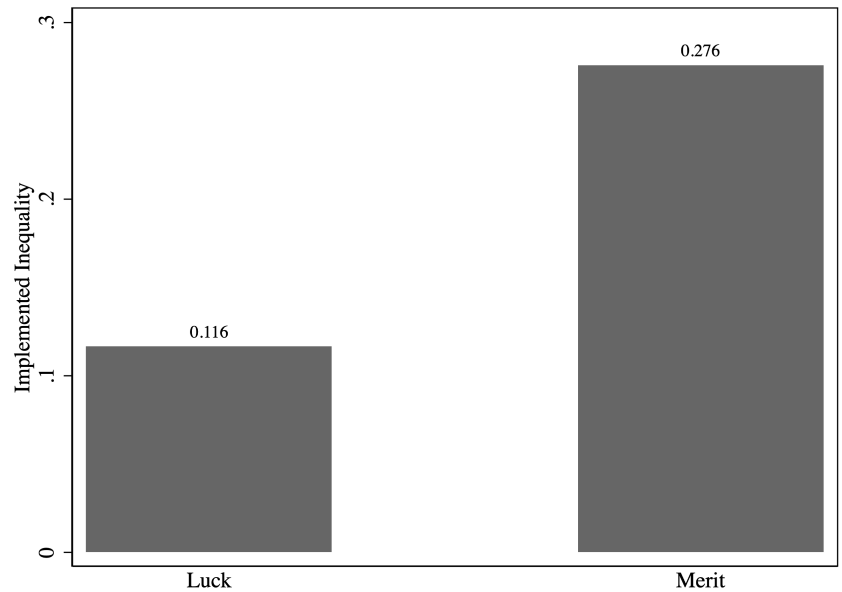

(b) Share of inequality that subjects choose to erase

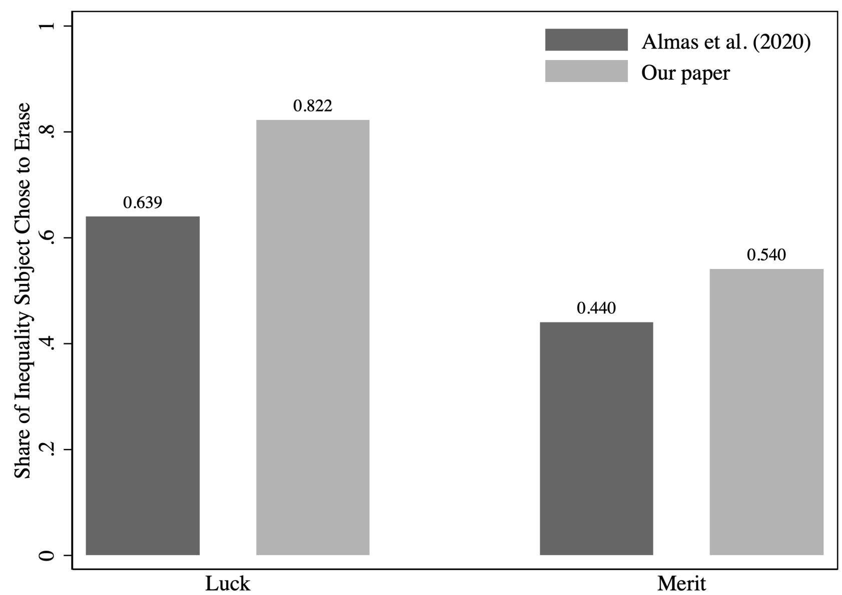

Notes: Figure (a) shows the average implemented inequality in treatment (merit) and control (luck) groups of the luck-vs-merit experiment in our paper following the measurement of Almås *et al.* (2020). Figure (b) shows the share of inequality that subject choose to erase in both Almås *et al.* (2020)'s and out paper. "Merit" group in the experiment means that tbe subject in that group was told that the hypothetical individuals had earned their endowments by their merits (answering SAT questions/productivity). "Luck" group includes those who were told their endowments were allocated based on a coin toss.

Figure C.3: The share of different fairness types

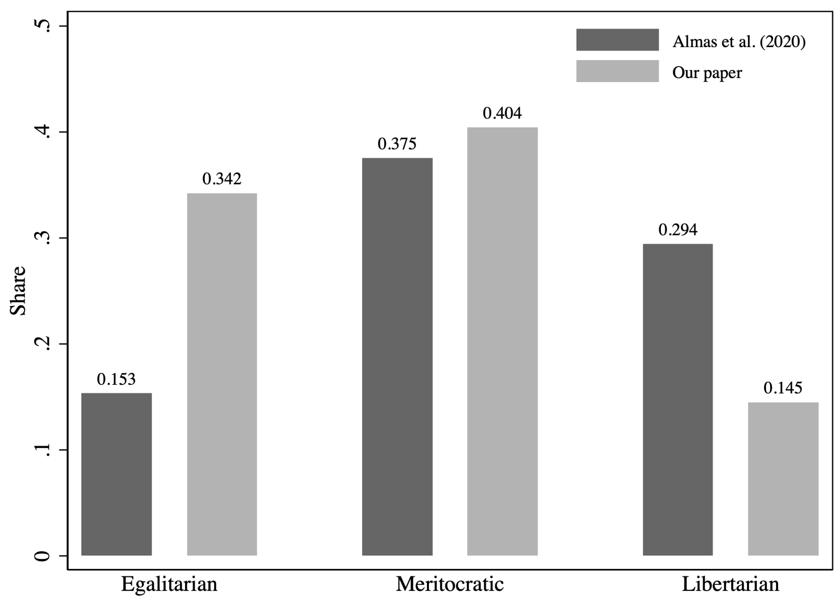

Notes: This figure shows the share of different fairness types in our luck-vs-merit sample.

D Data Appendix

Much of the data in this paper is collected via Amazon Mechanical Turk (MTurk), and in this short data appendix we detail our data collection process. MTurk allows "requesters" to find "workers" to complete "human intelligence tasks" (HITs). As of the summer of 2014 (when much of our data were collected), MTurk advertised that requesters could "access more than 500,000 workers."

We registered as a requester and posted the following HIT: "Short (less than ten minutes) opinion survey on a variety of topics." We tried to use a neutral description that would limit selection bias while also giving workers an honest description of the task. Compensation was set to \$1, approximately minimum wage assuming that subjects took seven minutes to complete it. Actual median completion time was 10.1 minutes, implying an hourly wage of \$6.09. Though we cannot find official data on average wages on MTurk, worker forum threads suggested that, at the time, we were paying a generous wage (and when we posted requests for 300 survey takers, the sample was generally gathered within an hour).

Each worker logs in with an MTurk ID. We collected MTurk data over eleven separate sessions, dropping workers who had taken a previous survey with the same ID, to ensure a fresh sample each time (though our main results hold when we keep repeat-takers in the sample, see Appendix Table D.1).[28]

An important concern with MTurk is the possibility of "bots," algorithms that masquerade as humans. We thus begin each survey with a "captcha" (non-standard writing difficult for computers to interpret). Furthermore, when we ask respondents for feedback at the end of the survey, essentially the only negative comment was that the captchas were "hard," suggesting that algorithms would have a difficult time parsing them. Nonetheless, to address worries that robots learn how to read "captchas," we hand-sketched pictures of a cat, dog, horse and panda bear, and respondents in later rounds were asked to answer multiple-choice questions of the form: "this is a picture of...." after seeing these sketches. If robots still remain in our sample after these checks, they would attenuate any treatment effect. Importantly, all of our MTurk data are collected well before 2018, when many researchers found that MTurk became infested with "bots."[29]

To limit heterogeneity of the sample, we collect all data on workdays during daylight hours on the U.S. East Coast. Individuals were automatically prompted for a response when they tried to skip questions (to discourage robots or inattentive respondents). Given our

[28]If workers have multiple IDs then some individuals may participate in multiple sessions. Outside of surveys (which appear to make up a very small share of all HITs), there would seem to be little financial incentive for MTurk workers to create multiple MTurk IDs, but we cannot completely eliminate the possibility that some have done so.

[29]See, e.g., Kennedy *et al.* (2018) and Dennis *et al.* (2018). These researchers refer to the summer of 2018 as the "MTurk quality crisis."

Table D.1: Main between-subject redistribution and tax results, includes subjects that participated in multiple batches

	Amount redistributed	Chosen tax rate
	(1)	(2)
Treated	-0.755***	-0.011**
	(0.092)	(0.006)
Cont. gp. mean	4.372	0.289
Observations	2763	1514

Notes: The first column replicates the analysis from col. (2) of Table 1 and the second replicates the analysis from col. (1) of Table A.6. "Treated" in col. (1) refers to the \$5/\$15 experiment (i.e., deciding how much to redistribute when the recipients know their initial endowments as opposed to when they do not, as in control condition) and "Treated" in col. (2) refers to the income tax survey experiment (i.e., deciding on ideal tax rate for person who received income increase five years ago as opposed to one year ago, as in control condition). $^{*}p < 0.1,^{**}p < 0.05,^{***}p < 0.01$.

focus in parts of the survey on American tax policy, we limited the survey's availability to those with U.S. billing addresses, and asked respondents to confirm their U.S. residency. To further ensure the attentiveness of our subjects, we limit respondents to those with positive ratings from at least ninety percent of past requesters.

E Screenshots of the Survey Questionnaire

Consider two other participants on mTurk, person A and person B. Based on a coin flip, we have given $5 to person A and $15 to person B.

You can now transfer money between persons A and B. *Persons A and B have already been told how much money we have given them*. If you decide to give Person A $X instead of $5, they will be told that they now have $X instead of $5. They will *not* know that there is another person (Person B) involved, or that a third party (you) determined the money they received.

Please indicate below what transfer, if any, you would make.

←Transfer more to Person A

Transfer more to Person B→

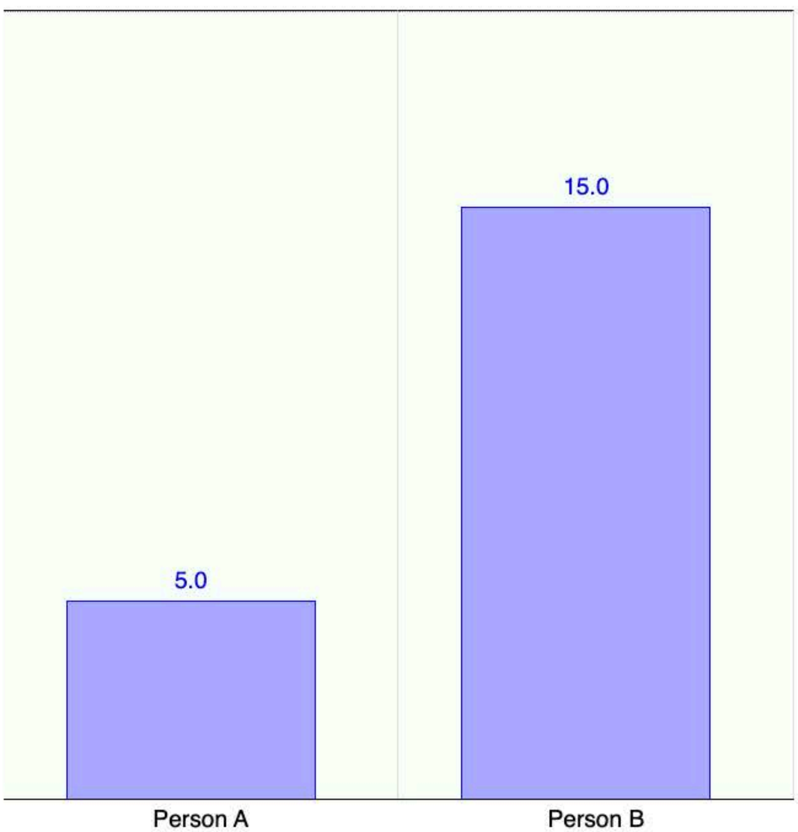

Recall: Person A and B have already been told how much money they were initially given.

$ 5.0 Person A

$ 15.0 Person B

>>

Survey Completion

0% 100%

Consider two other participants on mTurk, person C and person D. Based on a coin flip, we have given $5 to person C and $15 to person D.

You can now transfer money between persons C and D. *<u>Persons C and D are not told how much money they were initially given.</u>* If you decide to give Person C $X instead of $5, they will simply be told that they have been given $X, and will not know how much they started with. *Nor* will they know that there is another person (Person D) involved, or that a third party (you) determined the money they received.

Please indicate below what transfer, if any, you would make.

←Transfer more to Person C

Transfer more to Person D→

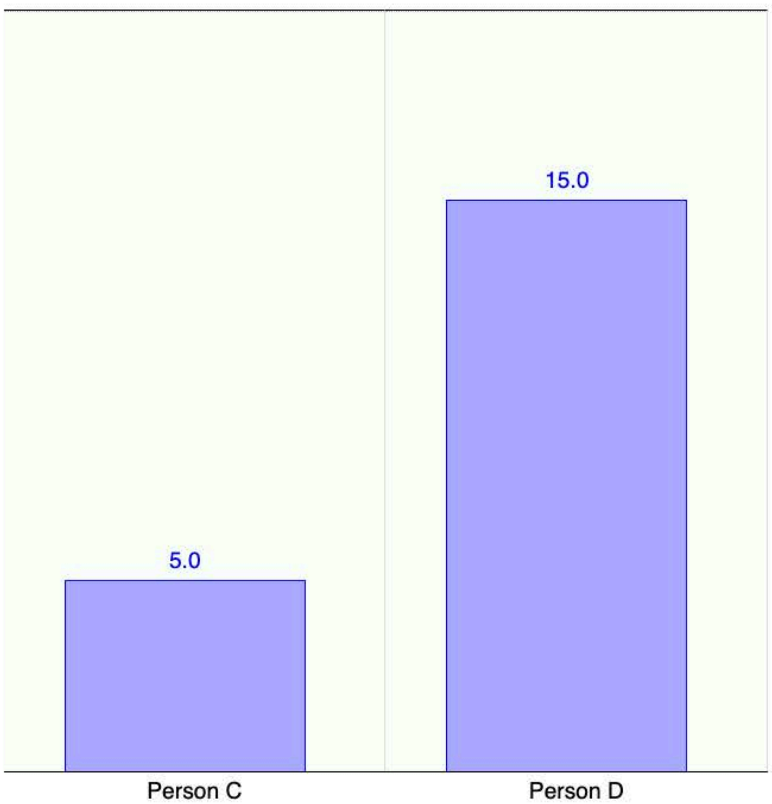

Recall: <u>Person C and D do NOT know how much money they were initially given.</u>

$ 5.0 Person C

$ 15.0 Person D

>>

Survey Completion

0% 100%

Suppose you bought a house for $250,000 a few years ago. The housing market in your neighborhood has since declined, and you have seen houses very similar to yours sell for $200,000, though some sell for a bit more and some sell for a bit less. You expect the current housing market conditions in your neighborhood to remain relatively stable. You are planning to relocate in the coming year for a new job. Someone is interested in buying your house. What is the least you would be willing to accept as a sale price? (Please only enter digits, i.e., no commas, decimals, or words.)

$ ______

Suppose you bought a house for $300,000 a few years ago. The housing market in your neighborhood has since declined, and you have seen houses very similar to yours sell for $200,000, though some sell for a bit more and some sell for a bit less. You expect the current housing market conditions in your neighborhood to remain relatively stable. You are planning to relocate in the coming year for a new job. Someone is interested in buying your house. What is the least you would be willing to accept as a sale price? (Please only enter digits, i.e., no commas, decimals, or words.)

$ ______

>>

Survey Completion

0% 100%

A small coffee shop has one employee who has worked there for six months and earns $10 per hour. The shop continues to do fairly good business, though unemployment in the area has increased due to a factory closure nearby. As a result, other small restaurants have now hired reliable workers at $8 an hour to perform jobs similar to those done by the coffee shop employee. The owner of the coffee shop reduces the employee's wage to $8. The owner’s actions were:

- ○ Completely fair
- ○ Acceptable
- ○ Unfair
- ○ Very unfair

A small coffee shop has one employee who has worked there for six months and earns $10 per hour. The shop continues to do fairly good business, though unemployment in the area has increased due to a factory closure nearby. As a result, other small restaurants have now hired reliable workers at $8 an hour to perform jobs similar to those done by the coffee shop employee. <u>The current employee leaves, and the owner decides to pay a replacement worker $8 an hour.</u> The owner’s actions were:

- ○ Completely fair
- ○ Acceptable
- ○ Unfair
- ○ Very unfair

>>

Survey Completion

0% 100%

Imagine that you face the following decisions. For each decision, please examine both options and indicate the one that you prefer.

Decision One:

- ○ A sure gain of $5.
- ○ 50% chance to gain $8 and 50% chance to gain nothing.

Decision Two:

- ○ A sure gain of $5.
- ○ 50% chance to gain $11 and 50% chance to gain nothing.

Decision Three:

- ○ A sure gain of $5.
- ○ 50% chance to gain $14 and 50% chance to gain nothing.

>>

Imagine that you face the following decisions. For each decision, please examine both options and indicate the one that you prefer.

Note that this time you will be making choices that involve *losses* of money.

Decision One:

- A sure loss of $5.
- 50% chance to lose $8 and 50% chance to lose nothing.

Decision Two:

- A sure loss of $5.
- 50% chance to lose $11 and 50% chance to lose nothing.

Decision Three:

- A sure loss of $5.
- 50% chance to lose $14 and 50% chance to lose nothing.

>>

The following statements inquire about your thoughts and feelings in a variety of situations. For each statement, indicate how well it describes you by choosing the appropriate letter on the scale under each question: A, B, C, D, or E. READ EACH ITEM CAREFULLY BEFORE RESPONDING. Please answer as honestly as you can.

I often have tender, concerned feelings for people less fortunate than me.

- ○ A (Does not describe me well)
- ○ B
- ○ C
- ○ D
- ○ E (Describes me very well)

I sometimes find it difficult to see things from the "other guy's" point of view.

- ○ A (Does not describe me well)
- ○ B
- ○ C
- ○ D
- ○ E (Describes me very well)

Sometimes I don't feel very sorry for other people when they are having problems.

- ○ A (Does not describe me well)
- ○ B
- ○ C
- ○ D
- ○ E (Describes me very well)

When I see someone being taken advantage of, I feel kind of protective towards them.

- A (Does not describe me well)
- B
- C
- D
- E (Describes me very well)

I sometimes try to understand my friends better by imagining how things look from their perspective.

- A (Does not describe me well)
- B
- C
- D
- E (Describes me very well)

Other people's misfortunes do not usually disturb me a great deal.

- A (Does not describe me well)
- B
- C
- D
- E (Describes me very well)

>>

Survey Completion

0% 100%

If I'm sure I'm right about something, I don't waste much time listening to other people's arguments.

- A (Does not describe me well)
- B
- C
- D
- E (Describes me very well)

After seeing a play or movie, I have felt as though I were one of the characters.

- A (Does not describe me well)
- B
- C
- D
- E (Describes me very well)

I believe that there are two sides to every question and try to look at them both.

- A (Does not describe me well)
- B
- C
- D
- E (Describes me very well)

I would describe myself as a pretty soft-hearted person.

- A (Does not describe me well)
- B
- C
- D
- E (Describes me very well)

When I'm upset at someone, I usually try to "put myself in his shoes" for a while.

- A (Does not describe me well)
- B
- C
- D
- E (Describes me very well)

>>

Survey Completion

0% 100%

There has been much talk about whether wealthy families are paying their fair share in taxes.

Consider the following person. He has been working for about five years as a regional sales manager at a medium-sized firm. *This year*, his firm was taken over by a larger corporation. While he will be doing the same job as before, to make his pay compatible with the earnings of employees in his position at the larger firm, his salary is now doubled, to $250,000.

If it were up to you, how much of his salary should he pay in taxes? (As a basis of comparison, the average American pays about 22 percent in taxes on the income they make.)

	0	10	20	30	40	50	60	70	80	90	100
Tax Rate (%)											

>>

Survey Completion

0% 100%

In general, how do you feel about increasing taxes on those making $250,000 or more (as has been proposed in Congress recently)?

- ○ Strongly oppose
- ○ Oppose
- ○ Favor
- ○ Strongly favor

>>

Survey Completion

0% 100%

Some people think that the government in Washington ought to reduce the income differences between the rich and the poor, perhaps by raising the taxes of wealthy families or by giving income assistance to the poor. Others think that the government should not concern itself with reducing this income difference between the rich and the poor. How do you feel? Think of a score of 1 as meaning that the government ought to reduce the income differences between rich and poor, and a score of 7 meaning that the government should not concern itself with reducing income differences. What score between 1 and 7 comes closest to the way you feel?

Please indicate your age.

Please indicate your gender.

- Male
- Female

>>

Survey Completion

0% 100%

What is your current employment status?

- ○ Full-time employee
- ○ Part-time employee
- ○ Self-employed or business owner
- ○ Unemployed and looking for work
- ○ Student
- ○ Not in labor force (for example: retired, or full-time parent)

What was your total household income, before taxes, last year (2012)?

- ○ $0-$9,999
- ○ $10,000-$14,999
- ○ $15,000-$19,999
- ○ $20,000-$29,999
- ○ $30,000-$39,999
- ○ $40,000-$49,999
- ○ $50,000-$74,999
- ○ $75,000-$99,999
- ○ $100,000-$124,999
- ○ $125,000-$149,999
- ○ $150,000-$199,999
- ○ $200,000+

Do you have children living with you?

- ○ Yes
- ○ No

>>

How would you describe your ethnicity/race?

- ○ European American/White
- ○ African American/Black
- ○ Hispanic/Latino
- ○ Asian/Asian American
- ○ Other

In which state do you live?

Please indicate your marital status.

- ○ Single
- ○ Married
- ○ Separated
- ○ Divorced

>>

Survey Completion

0% 100%

Which category best describes your highest level of education attained?

- ○ Eighth Grade or less
- ○ Some High School
- ○ High School degree/ GED
- ○ Some College
- ○ 2-year College Degree
- ○ 4-year College Degree
- ○ Doctoral Degree
- ○ Professional Degree (JD, MD, MBA)

Who did you support in the presidential election in 2012? If you were not able to vote, just choose the person you wanted to win the election at that time.

- ○ Barack Obama
- ○ Mitt Romney
- ○ Other

>>

Survey Completion

0% 100%

For the questions where you had to transfer money between other people, how did it change your thinking if these people had already been told how much money they were initially given?

Now we want your opinion about this survey. Your feedback is important to us. Did the survey feel biased toward a specific viewpoint?

- ○ Not biased
- ○ Biased toward a liberal viewpoint
- ○ Biased toward a conservative viewpoint
- ○ Biased toward another viewpoint

If you answered "biased toward another viewpoint," please explain.

Do you have any other thoughts about the survey that you would like to share?

Did you find any part of this survey confusing? If so, please explain.

Please provide your mTurk worker ID. We will use it to verify that you completed the survey.

>>

Survey Completion

0% 100%