Inequality of Opportunity, Biased Beliefs, and Demand for Redistribution: Pre-Analysis Plan

Marcel Preuss   Germán Reyes   Jason Somerville   Joy Wu

September 28, 2021

Contents

1 Introduction

Previous research has shown the majority of individuals in the United States hold a meritocratic fairness ideal (Almås et al., 2020). That is, most individuals accept income disparities that are due to differences in performance or effort but prefer to equalize inequalities that are mainly driven by pure luck or chance (Cappelen et al., 2007). However, the influence of luck on outcomes can come in different forms and at different points in time. First, luck may affect final outcomes in a way that is closely tied to the effort of an individual. A student may work hard to get accepted to college, but having the financial resources to attend college may be beyond her control. Alternatively, luck can sometimes affect outcomes independently of effort; for example, getting sick on the day of an important interview or exam can prevent someone from succeeding regardless of how well they were prepared. Second, luck may affect outcomes at different points in time. As in the example above, luck can materialize before or after exerting effort. For example, a student who does not know the outcome of a scholarship application will face uncertainty over the returns to studying. Conversely, learning about whether attending higher education is financially feasible in advance gives the student more control over the decision to pursue higher education.

We explore how both the nature and timing of luck affects individuals' views on inequality and demand for redistribution. Specifically, we ask if there are systematic differences in demand for redistribution when unequal earnings arise due to lucky opportunities or pure outcome luck. To answer this question, we implement a novel experimental design in which we hold fixed the impact of luck on outcomes while varying how luck manifests itself. To isolate the mechanisms driving any potential observed differences, we also experimentally vary the timing of when luck is realized and the transparency of the role of luck. The remainder of this pre-analysis plan presents the proposed experimental design, hypotheses, and analyses.

2 Research Strategy

We plan to collect data on redistribution choices in a sample that is representative of the U.S. population. A sample of impartial spectators will decide whether to adjust the initial earnings of workers. Workers' initial earnings will depend on some combination of effort and luck. To ensure that spectators' choices will actually affect workers' earnings, we will first recruit a sample of workers using a real online labor marketplace. We will then randomly pair these workers and assign an initial bonus payment of $5 to the winner. The spectators' task is to decide whether to redistribute earnings from the winner to the loser.

2.1 Recruitment of Workers

We will recruit participants on Amazon Mechanical Turk to participate in the worker task. We plan to recruit a sample of 2,400 workers. We require participants to be U.S.-based, have a 95% minimum approval rate, and have at least 500 approved HITs. We will pay workers a fixed participation fee of $2 upon completion, and they receive an additional payment of up $5 based on the decision of a randomly-chosen spectator.

2.2 Recruitment of Spectators

Will survey a pool of panelists that have recently participated in the *Survey of Consumer Expectations*. The survey will be administered by Nielsen, who will distribute direct links to our online experiment and manage participant payments. The target sample size is 1,200. We will anticipate two to three waves of recruitment over several weeks to reach our target sample size. The sample will contain current and recent panelists who have taken part in the *Survey of Consumer Expectations* over the past two years. Our experimental interface is mobile-friendly to encourage hard-to-reach demographic groups to participate in our experiment.

3 Design

In our experiment, we plan to take the first steps to empirically test whether observationally equivalent forms of luck that manifest through different channels lead to differences in implemented redistribution. Our experiment has two types of agents: workers and spectators. The experiment is divided into three stages: a production stage, an earnings stage, and a redistribution stage. In the production stage, workers engage in a real-effort task for a fixed amount of time. Payments are determined by randomly pairing workers and determining a winner and loser, which will depend on both worker performance and pure chance. In the earnings stage, winners are initially allocated earnings of $5, and losers allocated $0. In the redistribution stage, third-party spectators have the opportunity to redistribute earnings for a pair of real workers from the production stage. Spectators are always informed of the workers' outcomes (but not their outputs) and the knowledge workers had about the potential influence of luck and effort on their earnings. Moreover, spectators always have some information about whether luck was involved (design specific).

3.1 Worker Task

Workers complete an effort-based transcription task, transcribing as many 3-letter "words" into numerical code within 5 minutes. Performance is incentive-compatible towards increasing the chance of winning a tournament against a randomly-assigned, anonymous opponent. All workers are also subject to some form of exogenous luck that can influence their outcome against their opponent. Since the behavior of workers is not essential for our research question, we only provide the information necessary in relation to the spectator part and a brief discussion of worker exclusion criteria.

Worker Exclusion Criteria: The final set of workers *excludes* participants who fail to meet our pre-specified attention checks. After instructions are provided, the worker is tasked with answering five multiple-choice comprehension questions. The experiment records the total number of incorrect responses across all attempts. If the number of incorrect responses exceeds three, then the participant is automatically ineligible to continue, and a new participant is recruited. The survey cannot proceed unless all comprehension questions are answered correctly. If at least one response submitted is incorrect, then the page reloads with the respondent's previous selection saved but with flags indicating the incorrect responses.

We will also exclude from the analysis workers who do not complete any transcription tasks. Workers are encouraged to continuously work throughout the five minutes and make an "honest attempt". Workers who are inactivate for a continuous period of 90 or more seconds are excluded from our sample. Workers who meet these exclusion criteria are prevented from survey completion in real-time, and a new participant is recruited from the sample pool. Based on our pilot testing from the worker sample pool, these criteria are expected to exclude approximately 15–25% of workers in our sample.

3.2 Spectator Task

In the main experimental part of this study, a sample of third-party spectators makes earnings redistribution decisions between pairs of workers. In the earnings phase of the study, we randomly pair workers and determine the winner based on some combination of effort and luck (see Section 3.2.2 for group randomization). The winner receives a \$5 bonus payment, and the loser receives no additional payment. The spectators' task is to decide whether to redistribute earning between two randomly paired workers and, if so, by how much.

Spectators make a total of 12 redistribution decisions involving randomly paired workers, with each decision varying in the degree of luck involved in the worker-pair outcome. The precise nature of luck (i.e., group randomization) is described in detail in Section 3.2.2. Spectators are fully informed about what information the workers had available to them and when that information was revealed.

3.2.1 Redistribution Choices

The spectators' task is to choose how much income to redistribute from the winner to the loser. They can choose to redistribute any amount from \$0 to \$5 in \$0.5 increments. Each decision is presented in the form of an adjustment schedule. A \$0.00 adjustment (i.e., \$5.00 to the winner and \$0.00 to the loser) is the first option and is labeled as a "no"-adjustment choice. The remaining $\{\$0.50, \ldots, \$5.00\}$ redistribution choices are labeled as a "yes" redistribution choice, and denote the final earnings for both the winner and the loser: that is, $\{$(winner gets, loser gets)$\} = \{(\$4.50, \$0.50), (\$4.00, \$1.00), \ldots, (\$0.50, \$4.50), (\$0.00, \$5.00)\}$.

We provide four pieces of information to spectators on each redistribution decision screen.

Table 1: Information on Worker-Pairs Common to All Spectators

Decision X of 12		
Worker ID:	1bx64fef	1uj72mti
Result:	won	lost
Unadjusted Earnings:	$5.00	$0.00

First, we list worker IDs to emphasize further that these workers are real individuals. Second, we inform the spectators of the worker outcomes—that is, which worker won the tournament. Third, we provide the unadjusted earnings, which is always $5 for the winner and $0 for the loser. Finally, and most importantly, we provide information on the role luck played in determining the outcome. Specifically, we provide either the individual multipliers that each worker received or the probability that the outcome was determined by a coin flip (see Section 3.2.2 for details).

Spectators' responses are incentive-compatible, and one of their decisions will determine the final, adjusted earnings of a real pair of workers. We randomly select one of the 12 decisions and implement it for that pair of real workers. We also emphasize to spectators that they should treat each decision as if it will be implemented. Importantly, we also inform spectators that workers are not told whether they won or lost or the exact amount they will earn in each case. Spectators are told that workers were informed that they could earn up to $5 and that winning against their randomly assigned opponent increases their chances of earning more. This design intentionally removes any confounding issues relating to spectators' unwillingness to make workers know that money has been taken (or added) to an expected or reference amount. Spectators are also assured that while workers know a third party may influence their final earnings, the spectators' identity is completely anonymous to workers.

3.2.2 Spectator Treatments

Our experiment follows a 3×2 between-subjects design. First, spectators will be randomly assigned to one of three luck conditions:

- *CF* (Coin Flip): Redistribute earnings for pairs of workers who, with some probability q, the winner was determined by a fair coin flip rather than performance.
- *EA* (Ex Ante): Redistribute earnings for pairs of workers who are assigned a score multiplier, m_i, *before* beginning of their transcription task as a rate of return on their performance.
- *EP* (Ex Post): Redistribute earnings for pairs of workers who are assigned a score multiplier, m_i, *after* completing the transcription task as a rate of return on their performance.

In the EA and EP treatments, spectators know which player won and the multipliers of each player, (m_i, m_j), but not the number of tasks completed by each worker. Panel A of Table 2 provides an example of a typical trial from the EA or EP condition. In the CF treatments, spectators know the probability q that a coin flip determined the outcome. Spectators also know that this probability is not revealed to workers, despite workers knowing that there is some (unstated) chance a coin flip would determine their outcomes. Panel B of Table 2 provides an example of a typical trial from the CF condition.

Table 2: Example of spectator trials for each luck treatment

Panel A: Multiplier (EA and EP) trial		
Worker ID:	ga2c8k8x	nkqqjd0n
Multiplier:	2.9	2.4
Result:	won	lost
Unadjusted Earnings:	\$5.00	\$0.00
Panel B: Coin Flip (CF) trial		
Worker ID:	sao9rqhr	qeha27vh
Coin-Flip Chance:	46%	
Result:	won	lost
Unadjusted Earnings:	\$5.00	\$0.00

Panel A depicts an example multiplier (EA and EP) trial. Panel B depicts an example coin flip trial.

Our second between-subjects treatment is to randomly assign half of the spectators to receive info about the probability that the winner was the worker who completed more tasks.

- *INFO*: Provide information on the probability that the winner of a pair of workers was also the better performer.
- *NO INFO*: No further information provided.

In the CF-INFO group, spectators will see the following text: "This means that there is a π% chance that the winner above completed more transcriptions than the loser." In the EP-INFO and EA-INFO treatments, spectators will see the following text: "Based on historical data for these multipliers, there is a π% chance that the winner above completed more transcriptions than the loser".

The values of π vary from trial to trial within-subject. Spectators will make a total of 12 decisions. For each decision, spectators will see a different value of π. To select the 12 values of π, we constructed 12 bins:

$$\pi \in \Big\{\{50\}, \{51, ..., 54\}, \{55, 56, ..., 59\}, ..., \{95, 96, ..., 99\}, \{100\}\Big\}. \quad (1)$$

For each spectator, we randomly draw one value of π from each of the 12 bins. This process ensures that every spectator makes a decision with $\pi = 50$ and $\pi = 100$, and that the values of π observed by each spectator are distributed throughout the support of π. The order in which the 12 π values appear is randomized.

Mapping m's and q to π. The key information we present on each trial is the multiplier of each player, (m_i, m_j), or the probability that the winner was determined by a coin flip, q. It is therefore necessary to map each π value to corresponding (m_i, m_j) or a coin-flip probability q. Let π denote the probability that the winner was the person who completed more transcriptions.

Mapping π to q is straightforward and is given by the formula $\pi = 1 - 0.5q$. Mapping π to a (m_i, m_j) is less straightforward. In Section 4, we derive the formula for π for meritocrats. To estimate π for a given m, we can examine all possible worker pairs and then compute the fraction of times that the winner was the person who solved more translations when each of the workers is assigned a relative multiplier of m. With 800 workers per treatment, there are $(800(800-1))/2 = 319,600$ possible pairings. Since the higher multiplier can be assigned to either worker, that creates 639,200 observations to calculate π for each relative multiplier, m. Using this method, we can compute for any given m, how likely it is that the winner exerted more effort. This provides a mapping from m to π in the EA and EP treatments.

Sub-treatments. In previous studies, the rules of how the winner is determined are revealed after workers complete the task (e.g,. Cappelen et al. (2019)). In our EA condition, it is necessary to reveal these rules before workers begin working. For comparability with previous literature and to provide a clean comparison between the EA and EP/CF conditions, we randomly vary the timing of when the rules for determining the winner are revealed for the CF and EP conditions.

- *RULES*: Workers are informed that there will be multipliers (EA or EP) or a coin flip (CF) that influences the final outcomes *before* they start their task.
- *NO RULES*: Workers are informed that there will be multipliers (EP) or a coin flip (CF) that influences the final outcomes *after* they complete their task.[1]

Table 3 summarizes our between-subjects treatments and the fraction of our sample that will be assigned to each one.

[1] Note that this cannot be a condition for EA participants. Refer to * in Table 3.

Table 3: Between-Subjects spectator randomization groups

Treatment	**Randomization Group Label (% of Sample)**		
Luck Type	EA (33.3%)	EP (33.3%)	CF (33.3%)
Effort Info.	Info, No Info (16.7%, 16.7%)	Info, No Info (16.7%, 16.7%)	Info, No Info (16.7%, 16.7%)
Rules	Rules Only* {(16.7%, 16.7%)}	Rules, No Rules {(8.3%, 8.3%), (8.3%, 8.3%)}	Rules, No Rules {(8.3%,8.3%), (8.3%, 8.3%)}

3.2.3 Belief and Background Questions

After spectators complete the 12 redistribution decisions, we ask them a series of follow-up questions. These questions are designed to be exploratory and to permit some heterogeneity analyses. First, we present spectators with a randomly selected trial that they saw before. Instead of asking them to make a redistribution decision, we instead ask what they believe is the likelihood that the winner completed more transcriptions than the loser. Secondly, we ask them what they think the performance of the worker with the lower multiplier was. Thirdly, we include Likert-type survey questions related to the spectators' fairness beliefs related to income, effort and luck in the real world. One attention check is also included in this Likert-type series.

3.2.4 Pilot

We conducted two spectator pilot studies on Prolific to test the implementation logistics of our design and to help clarify certain design elements. We first ran a pilot with 50 participants in each treatment. We also ran a follow-up pilot with 50 participants in the EA treatment only. We used this pilot data to help calibrate the comprehension questions and to minimize survey fatigue. We also used timing data to set an appropriate participation fee for our main sample. This pilot was also used to evaluate the comprehension of the instructions.

4 Framework

In this section, we build a model of demand for redistribution in the presence of uncertainty about worker effort. The framework stresses the importance of beliefs for supporting redistribution and highlights possible sources of bias. The theoretical results from this section motivate and guide our main experimental hypotheses.

In the redistribution stage, spectators decide how much income to redistribute between workers. Let f be the fair-income share of the high-effort worker. We allow f to vary across treatments because different conditions may impose different perceived difficulty, stress, or other factors on workers.

We follow Cappelen et al. (2019) and assume that the fair-income share preferences are not a function of the difference in effort exerted by both players, but rather that it only depends on which player exerted more effort. The main hypotheses that we present in this section remain unchanged if we relax this assumption. However, making this simplifying assumption here helps to develop an intuition for our hypotheses in the most concise way.

Assuming a quadratic loss function, the problem of each spectator is to choose the fraction of earnings allocated to the winning worker, y, to maximize

$$\max_{y\in[0,1]} \quad -(y-f)^2. \tag{2}$$

Spectators observe the outcome of each match but do not directly observe the effort level of each worker. To maximize (2), the spectator has to infer the probability that worker i exerted the highest level of effort given that they had a higher score than worker j. Let π denote such a probability. Given π, the spectator's problem is to choose the fraction of earnings allocated to the winning worker, y, that maximizes the expected utility

$$\mathbb{E}(U(\cdot)) = -\pi\Big(y-f\Big)^2 - (1-\pi)\Big(y-(1-f)\Big)^2. \tag{3}$$

Then the optimal income share allocated to the winning worker in an interior solution is

$$y^*(\pi; f) = \pi f + (1-\pi)(1-f). \tag{4}$$

This equation highlights that redistribution depends on both preferences for redistribution *and* beliefs. The optimal amount of income for the winner increases linearly in π.

As is often the case in real life, spectators in our experiment do not observe π. Instead, they must form an estimate of π based on the other signals about the importance of luck. In the CF condition, the coin-flip probability q is a sufficient statistic to derive π. In the EA and EP conditions, spectators must form an estimate of π using information about the relative magnitudes of m_i and m_j and the perceived effort distribution. Thus, the role of luck is more transparent in the CF condition than in the EA and EP conditions.

More formally, to infer π from q, spectators must solve a relatively simple Bayesian updating problem: $\pi = 1 - 0.5q$. To assess π in the EA and EP conditions, spectators must correctly guess the distribution of effort first. Then, π follows from

$$\pi = \Pr_\kappa\Big(e_i \geq e_j \Big| m_i e_i > m_j e_j, m_i, m_j\Big), \tag{5}$$

where the subscript $\kappa \in \{EP, EA\}$ accounts for the possibility that the distribution of effort might vary across treatments. Without loss of generality, assume worker i had the highest score; that is, $m_i e_i > m_j e_j$. There are two cases that spectators must consider: $m_i \leq m_j$ or $m_i > m_j$. If $m_i \leq m_j$, then $\pi = 1$. Intuitively, if worker i won *despite* having a lower multiplier, then it must be the case that they exerted the highest effort level. Conversely, if $m_i > m_j$, equation (5) becomes

$$\pi = \frac{\Pr\left(e_i \geq e_j | m_i, m_j\right)}{\Pr\left(m_i e_i > m_j e_j | m_i, m_j\right)}. \tag{6}$$

Importantly, π depends only on the relative multiplier m. Moreover, π is convex and decreasing in m if the distribution of effort is bell-shaped; that is, if it is unimodal with approximately equal mean and mode. Intuitively, it is convex because even a small multiplier difference is likely to have a big impact on who wins if worker effort tends to be similar.

The spectator's estimate of π may not be accurate for several reasons. In the CF condition, spectators may fail to perform basic Bayesian updating. In the EP and EA conditions, spectators must form an estimate of the worker-performance distribution when trying to infer π and engage in Bayesian updating. As part of the latter, spectators must estimate $\Pr_\kappa\left(m_i e_i > m_j e_j | m_i, m_j\right)$. To guess this term correctly (on average), they must understand that small increases in the multiplier difference can cause a large increase in the likelihood of winning. This relates to the transparency of luck, which may be more obvious if luck is presented as a coin flip than as score multipliers. Thus, let $\tilde{\pi}$ be the spectator's subjective estimate of π. Note that each $\pi \in [0.5, 1)$, π corresponds to a unique relative multiplier m in the EA and EP conditions, and to a unique value of q in the CF conditions. Thus, instead of writing $\tilde{\pi}$ as a function of (m_i, m_j) or q, it can be expressed as a function of the treatment and the true π directly, i.e. $\tilde{\pi}^\tau(\pi)$. In summary, we expect spectators to redistribute:

$$y^*(\tilde{\pi}^\tau(\pi), f^\tau) = \tilde{\pi}^\tau(\pi) f^\tau + (1 - \tilde{\pi}^\tau(\pi))(1 - f^\tau), \tag{7}$$

where $\tau \in \{CF, EP, EA\}$ indicates that both the perceived likelihood that the winner was more productive, given by $\tilde{\pi}^\tau(\pi)$, and the fair share for the more productive worker, given by f^τ, may depend on the treatment. To simplify notation, we omit the variable π from $\tilde{\pi}$ when it does not cause confusion. Expression (7) forms the basis for our experimental hypotheses. It highlights that biased beliefs may conflate the role of preferences in driving redistribution.

5 Hypotheses and Empirical Strategy

5.1 Preliminary Hypotheses

First, we expect to replicate a key finding from the literature. In the CF treatment, redistribution is decreasing in q, the probability that the winner was determined by a coin flip (H0a); for example, see Cappelen et al. (2007). Our EA and EP treatments are novel because luck manifests itself through multipliers as opposed to coin flips. However, we expect to find a similar relationship between luck and redistribution. Specifically, we expect the level of redistribution to be decreasing in π in the EP (H0b) and EA (H0c) conditions.

5.2 Main Hypotheses

Our main hypotheses concern how redistribution decisions differ across our experimental treatments and π. The primary comparison of interest is how redistribution under inequality of opportunity (EA) differs from redistribution under simple outcome inequality (CF). There are two ways in which redistribution may be different across our treatments. First, holding π fixed, the absolute *level* of redistribution may differ. Second, the *slope* of the relationship between redistribution and π not be the same. Because our EA and CF treatments also differ in the information that workers had about their luck when they began working, we leverage our EP treatment to distinguish changes in redistribution that stem from the timing of information versus the nature of luck. Thus, we focus on the following comparisons:

1. EA (no info) vs. EP (no info)
2. EP (no info) vs. CF (no info)

Comparing our EA and EP treatments allows us to isolate how the timing of when luck occurs influences demand for redistribution.[2] Comparing our EP and CF treatments allows us to assess whether redistribution depends on the nature of how luck manifests itself. Our first test concerns whether the slopes differ between our treatments. Formally, we test the null that the slope of redistribution with respect to π is the same across our treatment comparisons:

$$\frac{\partial y^*(\tilde{\pi}^{EA}, f^{EA})}{\partial \pi} = \frac{\partial y^*(\tilde{\pi}^{EP}, f^{EP})}{\partial \pi} \text{ for a given } \pi \in [0.5, 1] \tag{H1-a}$$

$$\frac{\partial y^*(\tilde{\pi}^{EP}, f^{EP})}{\partial \pi} = \frac{\partial y^*(\tilde{\pi}^{CF}, f^{CF})}{\partial \pi} \text{ for a given } \pi \in [0.5, 1] \tag{H1-b}$$

In other words, we assess whether redistribution responds to changes in luck the same across our treatments.

Second, we assess whether the average level of redistribution differs between treatments for a given π. Rather than comparing average slopes across all π as in H1a to H1c, this allows for a more granular approach to assessing the shape of the relationship between redistribution and luck. Formally, we test the following nulls:

$$y^*(\tilde{\pi}^{EA}, f^{EA}) = y^*(\tilde{\pi}^{EP}, f^{EP}) \text{ for a given } \pi \in [0.5, 1] \tag{H2-a}$$

$$y^*(\tilde{\pi}^{EP}, f^{EP}) = y^*(\tilde{\pi}^{CF}, f^{CF}) \text{ for a given } \pi \in [0.5, 1] \tag{H2-b}$$

[2]Notice that rules are always announced to workers before beginning the task in the EA treatment. Hence, the clean comparison is EA (NO-INFO) vs. EP (NO-INFO and RULES). If we find that the timing of when the rules are revealed does not have a statistically significant effect on redistribution in the EP treatment, we will pool the EP rules treatments.

5.3 Controlling for Information

Our information treatment is designed to distinguish the role of preferences versus biased beliefs in driving any observed differences in demand for redistribution. By presenting spectators with accurate information about π, we eliminate the role of biased beliefs and enable those who care about π to make better distributive choices. By providing information about π, we can isolate the role of preferences in driving any observed differences in the slopes that we identify. Formally, if $\tilde{\pi}^\tau(\pi) \approx \pi$ in our information treatment, then $\partial y^*(\tilde{\pi}^\tau, f^\tau)/\partial\pi = 2f^\tau - 1$. Therefore, we test the following null hypotheses:

$$\frac{\partial y^*(\pi, f^{EA})}{\partial \pi} = \frac{\partial y^*(\pi, f^{EP})}{\partial \pi} \text{ for a given } \pi \in [0.5, 1] \tag{H3a}$$

$$\frac{\partial y^*(\pi, f^{EP})}{\partial \pi} = \frac{\partial y^*(\pi, f^{CF})}{\partial \pi} \text{ for a given } \pi \in [0.5, 1] \tag{H3b}$$

Note that we can identify whether spectators are over- or under-reacting to luck by comparing slopes across the INFO and NO-INFO treatments in each treatment. Formally, note that:

$$\frac{\dfrac{\partial y^*(\tilde{\pi}^\tau, f^\tau)}{\partial \pi}}{\dfrac{\partial y^*(\pi, f^\tau)}{\partial \pi}} = \frac{\partial \tilde{\pi}^\tau}{\partial \pi}$$

Thus, we test the following null hypotheses:

$$\frac{\partial \tilde{\pi}^{EA}}{\partial \pi} = \frac{\partial \tilde{\pi}^{EP}}{\partial \pi} \text{ for a given } \pi \in [0.5, 1] \tag{H4a}$$

$$\frac{\partial \tilde{\pi}^{EP}}{\partial \pi} = \frac{\partial \tilde{\pi}^{CF}}{\partial \pi} \text{ for a given } \pi \in [0.5, 1] \tag{H4b}$$

Our information intervention also allows us to identify the presence of biased beliefs at a given π. Provided $\tilde{\pi}^\tau \approx \pi$ in our information treatment, we expect redistribution in the information treatment to change by

$$y^*(\pi, f^\tau) - y^*(\tilde{\pi}, f^\tau) = (\pi - \tilde{\pi})(2f^\tau - 1). \tag{8}$$

As long as spectators are not pure egalitarians in all treatments ($f^\tau = 1/2$), we can test for the existence of biased beliefs by comparing redistribution in our INFO treatment to our NO-INFO treatments. Specifically, we test the following null hypotheses:

$$y^*(\pi, f^{EA}) = y^*(\tilde{\pi}^{EA}(\pi), f^{EA}) \text{ for any } \pi \in [0, 1] \tag{H5a}$$

$$y^*(\pi, f^{EP}) = y^*(\tilde{\pi}^{EP}(\pi), f^{EP}) \text{ for any } \pi \in [0, 1] \tag{H5b}$$

$$y^*(\pi, f^{CF}) = y^*(\tilde{\pi}^{CF}(\pi), f^{CF}) \text{ for any } \pi \in [0, 1]. \tag{H5c}$$

5.4 Heterogeneity

We will analyze heterogeneity in spectators' redistributive preferences based on gender, race, educational attainment, income, political orientation, attitudes towards redistribution and the government, and self-reported fairness ideals. We will conduct the heterogeneity analysis separately for each treatment condition as well as pooling across conditions.

5.5 Specifications and Analysis

Our empirical tests are based on a series of regression specifications. The outcome variable we focus on is the fraction of earnings that spectator i redistributes to the loser in worker pair p; that is, $y_{ip} \equiv 1 - y^*(\tilde{\pi}_{ip}, f_i)$. When $y_{ip} = 0$, the loser gets 0% of the total earnings and the winner retains 100%. If $y_{ip} = 0.5$, then both the winner and loser received half of the total earnings each. Our baseline specification relates the share of earnings redistributed by spectator to the ex-ante likelihood that the winner exerted more effort, π_{ip}:

$$y_{ip} = \alpha + \beta\pi_{ip} + \varepsilon_{ip}, \tag{9}$$

where ε_{ip} is an error term. In all regression models, We cluster standard errors at the spectator level.

Our empirical tests start by looking the average slope of redistribution with respect to luck. To test our preliminary (H0a-H0c) and main hypotheses (H1a−H3c), we will estimate a specification that allows the coefficients $\{\alpha_\tau, \beta_\tau\}$ to vary for each of our three treatment groups, $\tau \in$ {CF, EA, EP}. Specifically we will estimate the following equation using data from our NO INFO condition:

$$y_{ip} = \sum_{\tau} \tau_i \left(\alpha_\tau + \beta_\tau \pi_{ip}\right) + \varepsilon_{ip}, \tag{10}$$

where τ_i is equal to one if spectator i is in treatment τ and zero otherwise. The estimates from equation (10) allow us to assess a number of our hypotheses. First, we can test our preliminary hypotheses H0a to H0c by assessing whether $\hat{\beta}_{CF}, \hat{\beta}_{EA}, \hat{\beta}_{EP} \leq 0$. To assess our main hypotheses, we can compare the coefficient estimates across conditions. For hypotheses H1a and H1b, we test for equality of our slope coefficient on average: that is, whether $\hat{\beta}_{EA} = \hat{\beta}_{EP}$ and $\hat{\beta}_{EP} = \hat{\beta}_{CF}$, respectively.

Hypotheses H2a and H2b ask whether the level of redistribution is the same across luck treatments at a given π. In practice, comparing the average amount of redistribution across treatments for all 51 values of π would require a very large sample size. Alternatively, we can compare the average level of redistribution in each of the experimental π bins as defined in equation (1). Let $b \in \{1, ..., 12\}$ index the 12 experimental π bins. We will estimate the following regression specification:

$$y_{ib} = \sum_{\tau} \sum_{b} \gamma_{\tau b} \tau_i \pi_b + \varepsilon_{ib} \tag{11}$$

where π_b equals one if π belongs to bin b. Thus, we can test H2a and H2b for a given π bin by assessing whether $\widehat{\gamma}_{EA,b} = \widehat{\gamma}_{EP,b}$ and $\widehat{\gamma}_{EP,b} = \widehat{\gamma}_{CF,b}$. This approach will allow us to take a more flexible approach to mapping out how spectators' redistribution decision respond to luck. Specifically, we can assess whether the linear specification in (10) is accurate. Formally, we can assess whether adding π-bin dummies to (10) yields significant deviations from linearity. More generally, we can assess the shape of luck-distribution relationship visually by plotting the average level of redistribution for each π and treatment. We will also estimate best-fit curves for each treatment using local-linear regressions.

Next, we assess the impact of providing information to spectators about π. To test H3a-H3c, we re-estimate equation (10) using our full data and add interaction terms to denote our information conditions. Formally, we estimate

$$y_{ip} = \sum_{\tau} \tau_i \left(\alpha_\tau + \beta_\tau \pi_{ip} + \delta_\tau INFO_i + \phi_\tau INFO_i \pi_{ip} \right) + \varepsilon_{ip}, \tag{12}$$

where $INFO_i$ is an indicator equal to one if spectator i is in the information condition. To evaluate H3a and H3b, we again start by comparing the average slopes across our information conditions. Formally, we test whether $\widehat{\phi}_{EA} = \widehat{\phi}_{EP}$ and $\widehat{\phi}_{EP} = \widehat{\phi}_{CF}$, respectively.

Specification (12) also allows us to test hypotheses H4a and H4b. The slope of redistribution with respect to π in treatment τ is given β_τ for the NO-INFO treatment and $\beta_\tau + \phi_\tau$ in the INFO treatment. Therefore, we can test H4a by assessing whether $(\widehat{\beta}_{EA} + \widehat{\phi}_{EA})/\widehat{\beta}_{EA} = (\widehat{\beta}_{EP} + \widehat{\phi}_{EP})/\widehat{\beta}_{EP}$ and H4b by assessing whether $(\widehat{\beta}_{EP} + \widehat{\phi}_{EP})/\widehat{\beta}_{EP} = (\widehat{\beta}_{CF} + \widehat{\phi}_{CF})/\widehat{\beta}_{CF}$.

To test our final hypotheses, we extend equation (11) to include an indicator for our information treatment. Specifically, we estimate the following regression:

$$y_{ib} = \sum_{\tau} \sum_{b} \tau_i \pi_b (\gamma_{\tau b} + \rho_{\tau b} INFO_i) + \varepsilon_{ib}. \tag{13}$$

Therefore, testing $\widehat{\rho}_{EA,b} = 0$, $\widehat{\rho}_{EP,b} = 0$, and $\widehat{\rho}_{CF,b} = 0$ provide straightforward assessments of H5a, H5b, and H5c, respectively. Moreover, this allows to take a more flexible approach to assessing H3a and H3b by comparing the shape of the redistribution-luck relationship.

If we find evidence for biased beliefs, we will explore how those beliefs deviate from the objective π values. We will focus this analysis on the EA and EP conditions. Each redistribution decision p is defined by a worker pair $\{j, j'\}$ (the pair for which the spectator is deciding how to allocate earnings). Let m_j and $m_{j'}$ be the multipliers of workers j and j' and arrange each worker pair such that $m_j \geq m_{j'}$. For treatment conditions EA and EP, we will estimate regressions of the form:

$$y_{ip} = \alpha_\kappa + f(m_j, m_{j'}) + \beta_\kappa \pi_{ip} + \varepsilon_{ip}. \tag{14}$$

We will take a data-driven approach to assess which functional form $f(\cdot)$ provides the best fit to data. We will estimate (14) using a linear function, $f(m_j, m_{j'}) = m_j - m_{j'}$, the ratio between multipliers, $f(m_j, m_{j'}) = m_j/m_{j'}$, as well as polynomials of each of these functions. Statistically, we will assess whether $\hat{\beta}_\kappa$ remains negative and statistically different from zero as we control for alternative specifications of $f(\cdot)$. This will allow us to assess whether spectators are approximating π using some other heuristic or strategy, such as ironing.[3]

As a robustness check, we will test our hypotheses using the earnings inequality—as measured by the Gini coefficient of the final distribution of earnings among each worker pair—as an alternative, dependent variable.

5.6 Sample Exclusion Criteria

As described in Section 3.1, exclusion criteria for workers are programmed into the survey and participants are excluded for inattention in real time. Therefore, workers are not excluded after data collection. Furthermore, due to the quality of this sample of spectator participants, we do not have stringent exclusion criteria. While we collect data on comprehension question attempts and attention checks in the survey, we intend to use these to measure the quality of our sample.

References

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.

Cappelen, A. W., Hole, A. D., and Sørensen, Erikand Tungodden, B. (2007). The pluralism of fairness ideals: An experimental approach. *American Economic Review*, 97(3):818–827.

Cappelen, A. W., Mollerstrom, J., Reme, B.-A., and Tungodden, B. (2019). A meritocratic origin of egalitarian behavior.

[3]We will also collect self-reported measures of beliefs about $\tilde{\pi}$ (for one randomly-selected worker pair per spectator) and f. We will use these measures for exploratory analysis and to provide additional suggestive evidence to test our hypothesis.