# Inequality of Opportunity and Income Redistribution*

Marcel Preuss Germán Reyes Jason Somerville Joy Wu

March 2025

**Abstract**

Opportunities, such as access to education or family background, shape income inequality by influencing the chances of economic success. Unequal opportunities create uncertainty about whether success is merit- or luck-based. We examine how uncertainty impacts preferences for redistribution, comparing it to more transparent forms of luck. Using a U.S.-representative sample, we elicit redistribution decisions in two environments: when workers' earnings are influenced by luck directly ("lucky outcomes") or indirectly ("lucky opportunities"). Spectators redistribute less and are less sensitive to differences in luck under opportunities. Inferential challenges drive this gap, which implies that individuals are less meritocratic than previous findings suggest.

**JEL codes:** C91, D63

*Preuss: Cornell University, email, mp2222@cornell.edu; Reyes: Middlebury College and IZA, email, greyes@middlebury.edu; Somerville: Federal Reserve Bank of New York, email, jason.somerville@ny.frb.org; Wu: University of British Columbia, email, joy.wu@sauder.ubc.ca. For helpful comments and suggestions, we thank Peter Andre, Maxim Bakhtin, Hans-Peter Grüner, Ori Heffetz, Margaret Jodlowski, Muriel Niederle, Alex Rees-Jones, Ted O'Donoghue, Ryan Oprea, Matthew Rabin, and seminar participants at the briq Institute on Behavior and Inequality, the Economic and Social Research Institute, the VfS Annual Conference 2022, the 2022 Big Data in Economics at ZBW-Leibniz Information Centre for Economics, the 2022 North-American Economic Science Association Conference, the 2023 CESifo Area Conference on Public Economics, the ORG seminar at LMU Munich, and the 2024 Spring Behavioral and Experimental Economics Research Conference. The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of New York or the Federal Reserve System. The experiment reported in this paper was pre-registered in the AEA RCT Registry in October 2021, under the ID AEARCTR-0008147. Details on our experimental and survey items are available in our Supplementary Materials. The experiment was reviewed and granted an exemption by the Institutional Review Board at Cornell University. This work has been supported in part by a grant from the Cornell Center for Social Sciences. This paper was edited by Anna Dreber Almenberg.

## 1 Introduction

The design of redistribution policies—from progressive taxation to social welfare programs—depends on what citizens consider to be a fair distribution of income. Experimental work has found that most people hold meritocratic fairness attitudes: they tolerate income disparities that are merit-based but choose to redistribute when income differences result from circumstances beyond individuals' control, such as luck (e.g., Cappelen et al., 2007, 2013; Almås, Cappelen, and Tungodden, 2020). However, the widespread acceptance of meritocratic principles contrasts sharply with recent trends in income inequality in the United States, where support for redistribution has remained stagnant despite a rising impact of circumstances beyond people's control (Chetty et al., 2014; Ashok, Kuziemko, and Washington, 2015; Bozio et al., 2024).

We propose that this apparent disconnect between experimental findings and observed redistribution trends is partly due to limitations in how prior studies implemented luck in lab settings. Previous work has primarily focused on redistribution behavior in situations where luck arises in a simple and transparent way, making it easy to assess whether inequality is merit- or luck-based. By contrast, in a survey we conducted of U.S. households, we found that Americans identify unequal opportunities that merely affect someone's *chances of success* as the most critical source of luck in life.[1] For example, the main sources of luck identified by these U.S. households include the family one is born into and the education one receives. The impact of these opportunities on an individual's success is hard to assess without also observing that individual's effort. Thus, a fundamental feature of unequal opportunities is that they create uncertainty about the reasons for someone's success. This uncertainty creates inferential challenges that prior studies have not examined.

To illustrate the difficulty of implementing meritocratic fairness ideals when there are unequal opportunities, consider whether the success of Microsoft founder Bill Gates was due to merit or luck. On one hand, he was fortunate to attend one of the few high schools that offered unlimited access to a computer programming terminal. This is just one of many lucky opportunities in Gates' lifetime that were instrumental to his eventual success (Frank, 2016). However, he is also known for a fierce work ethic, famously stating that he "didn't believe in weekends; didn't believe in vacations." Both luck and effort were instrumental in his professional success, and assessing the exact impact of each factor poses a challenging inference problem.

We recruited 2,400 workers on an online labor market to perform a real-effort task and randomly paired them to compete for a fixed prize in a winner-takes-all environment. Then, we asked 1,170 individuals ("spectators") from the New York Fed's *Survey of Consumer Expectations*—a nationally representative survey of U.S. households—to choose the final earnings allocation for pairs of workers. We randomly assigned spectators to one of two luck environments. In the *lucky outcomes* environment, we selected the winner of each worker pair either by the number of tasks completed (merit) or by a coin flip (luck). In the *lucky opportunities* environment, the winner of each pair was

[1] Approximately 1,000 panelists from the New York Fed's *Survey of Consumer Expectations* participated in this survey. We provide a detailed description and analysis of the survey in Appendix B.1.

the worker with the higher number of tasks completed (merit) multiplied by a randomly assigned productivity multiplier (luck). These productivity multipliers capture, in reduced form, the type of advantages that create uncertainty about the source of success: circumstances beyond individual control such as family background, educational quality, and early-life conditions. We generate exogenous variation in the probability that inequality is merit-based—that is, the probability that the winner of a pair is the worker who completed more tasks—by varying the coin-flip probability (in lucky outcomes) or the relative productivity multipliers (in lucky opportunities). We use this probability to compare redistribution across the different luck environments while holding luck (the likelihood that the winner of a pair is the worker who performed better) constant.

In our baseline treatment, spectators are not informed about the probability that inequality is merit-based. Instead, spectators observe information about the environment that conveys the importance of merit and luck in determining the winner. The inferential challenge they face is using this information to infer how likely it is that the inequality between workers is merit-based. Under *lucky outcomes*, spectators observe the probability that the winner was determined by a coin flip. This feature of the lucky outcomes environment is linearly related to the probability that inequality is merit-based. Under *lucky opportunities*, they observe the productivity multiplier of each worker. Mapping this feature of the lucky opportunities environment to the probability that inequality is merit-based is more challenging since a small change in the relative advantage of the winner can imply large changes in the probability that the winner is the worker who completed more tasks.[2]

Our main result is that spectators redistribute less when inequality arises from lucky opportunities compared to lucky outcomes. On average, spectators redistributed 27.6 percent of earnings from winner to loser under lucky outcomes but only 23.3 percent under lucky opportunities. This redistribution gap emerges only when there is uncertainty about whether inequality is merit-based. Specifically, when performance alone determines the winner (i.e., under equal productivity multipliers or zero probability of a coin flip), spectators redistribute similarly across environments. As the likelihood of luck-based inequality increases, redistribution responds strongly under lucky outcomes but remains more muted under lucky opportunities, creating a redistribution gap between environments. We estimate that a ten percentage point increase in the probability of luck-based inequality leads to a 3.7 percentage point increase in redistribution under lucky outcomes but only a 2.0 percentage point increase under lucky opportunities.

We introduce two additional treatments to investigate the mechanisms underlying the differences in redistribution across luck environments. First, to further examine the role of inferential challenges, we provide direct information about the likelihood that inequality is merit-based. Consistent with spectators facing inferential challenges, we find that this information significantly reduces the redistribution gap between luck environments. We show that the information acts by changing

[2]The inferential challenges spectators face are not unique to our experimental design, but an inherent feature of many environments in which workers have unequal opportunities. Specifically, the probability that the outcome is merit-based is decreasing and convex in the winning worker's relative advantage. In Online Appendix Sections B.2 and B.3, we show that the same relationship prevails even if lucky opportunities are not multiplicative.

the way in which spectators incorporate unequal opportunities into their redistribution decisions. Second, to isolate the role of potential effort responses in driving our results, we implement an "ex-post" lucky opportunities condition in which workers learn about their multiplier only *after* completing the task. We find no significant difference in redistribution between the baseline and ex-post conditions, indicating that effort responses do not explain the redistribution gap.

We primarily contribute to a broad literature that studies how the source of inequality affects redistribution. Numerous experimental studies have demonstrated that people tolerate more inequality when earnings differences stem from differences in effort or individual choices, rather than luck (Cappelen, Sørensen, and Tungodden, 2010; Cappelen et al., 2013; Durante, Putterman, and van der Weele, 2014; Mollerstrom, Reme, and Sørensen, 2015; Cappelen et al., 2022a; Almås, Cappelen, and Tungodden, 2020; Cappelen et al., 2022b).[3] We show that whether luck arises through unequal opportunities or unequal outcomes significantly affects how much inequality individuals tolerate. We also advance the methodology of redistribution experiments by introducing a common scale for quantifying the importance of luck under uncertainty. This scale provides a portable definition of luck that can be widely applied to different tournament environments.[4]

Our work also contributes to the emerging literature examining how individuals make redistribution decisions when luck stems from differences in workers' circumstances (Bhattacharya and Mollerstrom, 2022; Dong, Huang, and Lien, 2024; Andre, 2024). Andre (2024) investigates whether spectators hold workers responsible for the unequal incentives they face. He finds that disparities in randomly assigned piece-rate wages produce large differences in worker performance. However, spectators reward workers according to their effort and irrespective of how differences in piece rates impact workers' performance. Dong, Huang, and Lien (2024) examine redistribution when luck arises through different employment offers or training quality. They find that spectators only partially account for these differences in advantages when making redistribution decisions. Similarly, Bhattacharya and Mollerstrom (2022) examine redistribution decisions when luck determines whether individuals can work at all. They find that spectators tolerate more inequality when luck determines who is allowed to work than when luck directly determines outcomes.

Our paper advances this literature in two fundamental ways. First, we study an environment where the impact of unequal opportunities is uncertain, mirroring real-world settings where the relationship between advantages and success is ambiguous. This design reveals that inferential challenges—the difficulty in assessing how much luck versus merit drives inequality—are key to understanding redistribution decisions. Second, while prior work focuses on how spectators reward

[3]Other research on the factors influencing redistribution has examined the role of other-regarding preferences (e.g., Charness and Rabin, 2002), fairness ideals (e.g., Konow, 2000; Cappelen et al., 2007) and perceptions regarding mobility and income inequality (e.g., Cruces, Perez-Truglia, and Tetaz, 2013; Kuziemko et al., 2015; Alesina, Stantcheva, and Teso, 2018; Fehr, Müller, and Preuss, 2024).

[4]Our approach builds on Cappelen et al. (2022b), who study redistribution under lucky outcomes using variation in the probability that inequality is merit-based. We extend their framework by deriving and implementing an analogous measure for opportunity-based luck, enabling us to systematically compare how redistribution responds to different forms of luck while holding the magnitude of luck constant.

performance differences that arise from unequal opportunities, we examine situations where equally hard-working individuals can achieve vastly different outcomes purely due to their circumstances. This shows that spectators have different attitudes towards redistribution even when there are no performance differences across workers.

Finally, we contribute to the literature that examines heuristics and biases in the inference process by documenting their consequences in an important economic context. We document that in the lucky opportunities environment, redistribution is close to linear on the workers' multiplier difference. This "linearization heuristic" leads spectators to underestimate the impact of unequal opportunities and consequently to redistribute less than they would under lucky outcomes. This finding is consistent with previous work showing that individuals often struggle to estimate non-linear relationships (Larrick and Soll, 2008; Levy and Tasoff, 2016; Rees-Jones and Taubinsky, 2020), and fail to solve even simple Bayesian updating problems (Benjamin, 2019), particularly in environments with uncertainty that require contingent reasoning (Martínez-Marquina, Niederle, and Vespa, 2019). The results from our information provision treatment are consistent with work showing that drawing attention to an attribute increases individuals' sensitivity to that attribute (Conlon, 2025) and that people focus on readily available information (Enke, 2020; Graeber, 2023).

## 2 Theoretical Framework

This section presents a stylized model of spectators' redistribution decisions when spectators care about the source of inequality but worker performance is unobservable. Our aim is not to test a specific model of spectator behavior, but rather to use the model to clarify our experimental setup and provide a common framework for quantifying the impact of luck on outcomes, regardless of its source. The model setup closely follows that of Cappelen et al. (2022b) and applies the framework to different luck environments.

Consider an impartial spectator who observes initial earnings in a winner-takes-all environment in which two randomly paired workers compete for a fixed prize. Spectator $i$'s task is to choose $r_i$, the fraction of income to redistribute from the winner to the loser. We model $r_i$ as a function of fairness views and beliefs about how likely it is that outcomes are based on merit. Formally, let $f_i$ denote the share of total income for the worker with the lower performance that spectator $i$ deems to be fair, and let $1 - f_i$ denote the fair share for the higher performance worker. We assume that spectators choose $r_i$ to minimize the quadratic difference between the fair allocation $(f_i, 1 - f_i)$ and the actual allocation $(r_i, 1 - r_i)$:

$$U(r_i, f_i) = -(r_i - f_i)^2. \tag{1}$$

If spectators know with certainty that the winner is the worker who performed better, then they implement the fair allocation, $r_i^* = f_i$. However, in our experiment, as is often the case in reality, spectators do not observe workers' performance. This creates uncertainty about whether

the best performer won or lost. As a result, spectators maximize the expected utility

$$\mathbb{E}(U(r_i, f_i)) = -\pi\Big(r_i - f_i\Big)^2 - (1-\pi)\Big(r_i - (1-f_i)\Big)^2, \tag{2}$$

where $\pi$ denotes the probability that the best performer won, and $1-\pi$ is the probability that the best performer lost. In other words, $\pi$ represents the probability that the outcome is merit-based, while $1-\pi$ represents the probability that the outcome is luck-based. Thus, we often refer to $1-\pi$ as our measure of the degree of luck in the environment, which we use to compare redistribution choices across different luck conditions in our experiment.

In an interior solution to (2), the optimal level of redistribution is

$$r_i^* = \pi f_i + (1-\pi)(1-f_i). \tag{3}$$

Equation (3) highlights that redistribution depends on both the spectator preferences about the fair share of earnings for the lower- and higher-effort worker ($f_i$) and how likely it is that the outcome was due to luck ($1-\pi$). Provided the fair share for the lower performance worker satisfies $f_i < 1/2$, the optimal level of redistribution increases in $1-\pi$. In other words, the more likely it is that the better-performing worker lost, the more spectators redistribute. When performance solely determines the winner ($1-\pi = 0$), spectators redistribute the fair share to the loser, $r_i^* = f_i$. When outcomes are due to pure luck ($1-\pi = 1/2$), spectators equalize earnings and choose $r_i^* = 1/2$.

### 2.1 Outcome Luck and Opportunity Luck

We consider two work environments with different processes for generating worker earnings. In the *lucky outcomes* environment, the winner of each worker pair is determined by either luck *or* performance. Specifically, we assume that with some probability $q \in [0,1]$, a coin flip determines the winner of the match. Conversely, with probability $1-q$ the winner is the worker who exerted more effort.

In the *lucky opportunities* environment, each worker is assigned an individual-specific productivity multiplier, $m_i$, which influences their earnings. We assume that the distribution of productivity multipliers is exogenous to the distribution of effort, meaning that a worker's productivity is determined by factors outside of their control, such as innate abilities or acquired skills (e.g., Mirrlees, 1971).[5] The winner of each worker pair is determined by comparing their final output, which is calculated by multiplying each worker's productivity multiplier by their performance level.

We assume that spectators observe the features of the environment that determine whether the outcome is merit-based $\pi$, but do not observe $\pi$ directly. In the lucky outcomes environment, spectators observe the coin-flip probability $q$. To infer $\pi$ from $q$, spectators must use Bayesian

[5]This approach captures the idea that some workers may be more fortunate than others due to their inherent talents or the opportunities they have had to develop their skills, leading to higher earnings even when putting in the same level of effort as their peers.

updating, which implies $\pi = 1-q/2$. In the lucky opportunities environment, spectators observe the productivity multipliers of each worker. Spectators can use this information to estimate $\pi$. Without loss of generality, assume that worker 1 wins the tournament, which means $m_1 e_1 > m_2 e_2$. The probability that the worker who won is the worker who exerted more effort given the information observed by the spectator is

$$\pi = \Pr\Big(e_1 \geq e_2 \Big| m_1 e_1 > m_2 e_2, m_1, m_2\Big). \tag{4}$$

When inferring $\pi$ from the features of the environment, spectators can encounter cases with or without uncertainty about whether the best performer lost. First, if $m_1 \leq m_2$, then $\pi = 1$ and there is no uncertainty. Intuitively, if worker 1 wins despite being less or equally productive as worker 2, then they must have exerted more effort than worker 2. Second, if $m_1 > m_2$, there is uncertainty and equation (4) becomes

$$\pi = \frac{\Pr\Big(e_1 \geq e_2, m_1, m_2\Big)}{\Pr\Big(m_1 e_1 > m_2 e_2, m_1, m_2\Big)} = \frac{1/2}{\Pr\left(e_2/e_1 < m_1/m_2\right)} \geq \frac{1}{2}, \tag{5}$$

where the equality follows under the assumption that the distribution of effort is independent of the multiplier.[6] Expression (5) shows that $\pi$ depends on the relative multiplier $m_1/m_2$ or, equivalently, on the relative advantage conferred to worker 1. We show in Appendix B.2 that $\pi$ is decreasing and convex in this relative advantage for any log-normal distribution of effort. Thus, the probability that the best worker lost $(1-\pi)$ or, equivalently, the probability that the winner was luck-based, is increasing and concave in the winning worker's relative advantage. Intuitively, this concavity arises because the density of $e_2/e_1$, which measures the relative effort of the disadvantaged worker compared to the advantaged worker, is centered around one (equal effort) and decreases for larger values of $e_2/e_1$. Therefore, even small relative advantages can greatly impact who wins, whereas the incremental effect of increasing this relative advantage matters less.[7]

This framework illustrates how differences in spectators' beliefs about $\pi$ can lead to differences in redistribution when luck affects workers indirectly through unequal opportunities versus when luck can affect workers' earnings directly. A reason for potential variation in beliefs about $\pi$ across our luck environments is that estimating $\pi$ may be more challenging under lucky opportunities, particularly if spectators fail to appreciate that even a small multiplier advantage can correspond to a considerable probability that the winning worker is not the one with the higher performance.

[6] This assumption is not essential for our analysis but simplifies the exposition of the theory. We also show in Section 5.2 that this assumption is consistent with our experimental data.

[7] The convexity of $\pi$ (concavity of $1-\pi$) in the relative advantage is not unique to the way we chose to model unequal opportunities. We show in Appendix B.3 that implementing opportunity luck as an additive headstart also leads to a convex mapping from lucky headstarts to the likelihood that the winning worker exerted more effort ($\pi$) for the empirical distribution of worker effort that we observe in the experiment. This is because $e_2 - e_1$, which is the relevant measure of relative effort in this case, has the most probability mass near zero (equal effort), and therefore, the convexity intuition above still applies (see the last paragraph in Section B.2 for the complete theoretical argument). We also show in Appendix B.2 that the same intuition readily extends to other common effort distributions.

## 3 Experimental Design

### 3.1 Production, Earnings, and Redistribution Stages

The experiment follows the impartial-spectator paradigm in Cappelen et al. (2013). It consists of three stages: a production stage, an earnings stage, and a redistribution stage. In the production stage, workers complete a real-effort task to encrypt three-letter "words" into numerical code (Erkal, Gangadharan, and Nikiforakis, 2011). Workers have five minutes to correctly encrypt as many words as possible using a dynamic and randomly generated codebook for each word (Benndorf, Rau, and Sölch, 2019). In the earnings stage, we randomly pair workers and determine the winner based on some combination of performance and luck. We initially allocate earnings of \$5 to the winner and \$0 to the losers. The exact interaction between luck and performance forms our main experimental treatments, which we describe in Section 3.2.[8] Workers know that a third party may influence their final earnings, but the spectator's identity is entirely anonymous to them.[9]

In the redistribution stage, spectators choose how much income to redistribute from the winner to the loser. Spectators make 12 redistribution decisions involving different real pairs of workers. To incentivize spectators to respond truthfully, we randomly select and implement one of their 12 decisions. Spectators may redistribute any amount between \$0 and \$5 in \$0.50 increments at no cost.[10] We present each decision as an adjustment schedule (see Appendix Figure A1 for an example of a redistribution choice). For example, an adjustment of \$1.50 implies \$3.50 for the winner and \$1.50 for the loser. The first option is always a \$0.00 adjustment, which we label a "no"-adjustment choice. The remaining $\{\$0.50, \ldots, \$5.00\}$ redistribution choices are labeled as a "yes"-adjustment choice and denote the final earnings for both the winner and the loser: that is, $\{(\text{winner gets, loser gets})\} = \{(\$4.50, \$0.50), (\$4.00, \$1.00), \ldots, (\$0.50, \$4.50), (\$0.00, \$5.00)\}$.

### 3.2 Spectator Treatments

We implement within-subject variation in how likely it was that the winner was merit-based (i.e., $\pi$). Spectators always observe features of the environment that are informative about the probability that the inequality between workers is merit-based. Moreover, the experiment embeds between-subject variation in the source of luck (lucky opportunities vs. lucky outcomes), the timing of when luck is revealed to the workers (ex-post vs. ex-ante), the information available to spectators about the importance of luck (full vs. partial), and the timing of workers' awareness about the

[8] We designed all experimental programs in oTree (Chen, Schonger, and Wickens, 2016). Screenshots of our experimental design and procedure are available in the Supplementary Materials.

[9] To combat the influence of anchoring effects in the spectators' redistribution choices, we inform spectators that we did not tell workers whether they won or lost nor the exact amount they will earn in each case. Spectators know that we only informed workers that they could earn up to \$5 and that winning against their randomly assigned opponent increases their chances of earning those \$5. This design removes any confounding issues relating to spectators' unwillingness to take earnings away from what workers might expect.

[10] We implemented costless redistribution to preserve comparability with prior experimental work on redistributive behavior (e.g., Cappelen et al., 2013; Almås, Cappelen, and Tungodden, 2020). This approach also allows us to isolate how spectators respond to varying sources of luck.

source of luck (after vs. before). Appendix Table A1 provides an overview of these between-subject treatments.

**3.2.1 Lucky Outcomes vs. Lucky Opportunities.** We randomly assign one-third of the spectators to redistribute earnings under lucky outcomes and two-thirds to redistribute earnings under lucky opportunities. In the lucky outcomes condition, we select the winner based on a coin flip with probability $q$ and based on the higher number of correct encryptions with probability $1 - q$. In the lucky opportunities condition, we generate inequality of opportunity by randomly assigning productivity multipliers to workers. For example, a worker with a multiplier of 1.2 who solved 20 encryptions would have a score of 24, while a worker with a multiplier of 3.0 who solved ten encryptions would have a score of 30. The winner in each pair is the worker who has the higher score. We draw the multiplier for each worker $i$ from the distribution: $m_i = 1$ with probability 0.05, $m_i = 4$ with probability 0.05, and $m_i \sim U(1, 4)$ with probability 0.9.[11] We round all multipliers to the nearest tenth.

We do not inform spectators about the workers' actual performance, though we do provide some information about how likely it is that luck determined the winner. In the lucky outcomes condition, spectators know the probability $q$ that we determine the winner by a coin flip, but not whether a coin flip actually determined the winner. Spectators also know that we do not reveal this probability to workers, though workers know that there is some unstated chance that a coin flip determines their outcomes. In the lucky opportunities condition, spectators know each worker's multiplier. We inform spectators that workers only know of their own multiplier and the distribution from which it is drawn, but do not know anything about the worker they compete against. Appendix Figure A1 provides an example of a redistribution decision in our lucky outcomes condition and an example of a redistribution decision in our lucky opportunities condition.

**3.2.2 Varying the Importance of Luck.** We implement within-subject variation in how likely it is that the better-performing worker won across worker pairs. We implement this variation in the lucky outcomes environment by changing $q$ across matches. We implement this variation in the lucky opportunities environment by varying the relative workers' multipliers across worker pairs. We control for the importance of luck by introducing a portable common metric across environments: the probability that the winner in a given pair completed more encryptions ($\pi$). When $\pi = 0.50$, there was a 50 percent chance that the winner of the match was the one who performed better; for example, when a coin flip determined the outcome ($q = 1$) in the lucky outcomes environment or when the ratio between worker multipliers is sufficiently large (so that the worker with the high multiplier always won the match) in the lucky opportunities environment. When $\pi = 1$, there was a 100 percent chance that the match's winner performed better; for example, when $q = 0$ or when the winning worker had a weakly lower productivity multiplier than the losing worker.

[11] We chose a relatively high upper bound of $m = 4$ to ensure that it is possible to generate $\pi = 1/2$ in our lucky opportunities treatment ($\pi = 1/2$ when the advantaged worker wins always; as if pure chance determines the winner).

Spectators make redistribution decisions for a total of 12 worker pairs. Each worker pair corresponds to a unique value of $\pi$ drawn from one of the following 12 bins:

$$\pi \in \Big\{ \underbrace{\{0.50\}}_{\text{Bin 1}}, \underbrace{\{0.51, ..., 0.54\}}_{\text{Bin 2}}, \underbrace{\{0.55, 0.56, ..., 0.59\}}_{\text{Bin 3}}, ..., \underbrace{\{0.95, 0.96, ..., 0.99\}}_{\text{Bin 11}}, \underbrace{\{1\}}_{\text{Bin 12}} \Big\}. \tag{6}$$

For each spectator, we randomly draw one value of $\pi$ from each of the 12 bins. This ensures that every spectator makes a decision with $\pi = 0.5$, $\pi = 1$, and that the remaining values are evenly distributed throughout the support of $\pi$. We present the 12 decisions in random order.

The key information we present on each decision is the multiplier of each worker pair, $(m_i, m_j)$, or the ex-ante probability that a coin flip $q$ determined the winner. Every possible pair $(m_i, m_j)$ and coin-flip probability $q$ correspond to a unique $\pi$ value. The mapping from $\pi$ to $q$ is given by the formula $q = 2(1 - \pi)$. To map $\pi$ to a multiplier pair, $(m_i, m_j)$, it is sufficient to consider the relative multiplier $m \equiv \max\{m_i, m_j\}/\min\{m_i, m_j\}$. Given any relative multiplier $m$, we examine all possible worker pairs and compute the fraction of times that the winner was the worker who solved more encryptions, which is the empirical counterpart to the probability that the winner is the worker who performed better, $\pi$.[12] This method yields a mapping from $m$ to $\pi$ as depicted in Panel B in Appendix Figure A2, which shows that $\pi$ converges to 1/2 for all $m > 3$. For a given $\pi$, we then select a random worker pairing with a corresponding relative multiplier.

**3.2.3 Timing of Opportunity Luck.** We also randomly vary the timing of when luck occurs. In our baseline lucky opportunities condition, we inform workers of their multipliers before they begin working on the encryption task. In the ex-post lucky opportunities condition, workers learn their multipliers only after completing the task. We assigned half of the spectators in the lucky opportunities conditions to the baseline treatment and the other half to the ex-post treatment.

**3.2.4 Information Intervention.** We randomly assign half of the spectators in each treatment to receive full information about $\pi$. In the lucky opportunities condition, we present the following additional text on the redistribution decision screen: "Based on historical data for these multipliers, there is a $[\pi \times 100]\%$ chance that the winner above completed more transcriptions than the loser." In the lucky outcomes condition, the equivalent text is: "There is a $[\pi \times 100]\%$ chance that the winner above completed more encryptions than the loser." As noted above, the value of $\pi$ varies from decision to decision. Appendix Figure A1 shows an example of the decision screens for the information treatments for both luck environments.

[12] With 800 workers per condition, there are $\binom{800}{2} = 319,600$ possible pairings. Since we can assign the higher multiplier to either worker, that creates 639,200 observations that we can use to calculate $\pi$ for each relative multiplier, $m$. Using this method, we compute, for each $m$, the fraction of all possible pairings in which the winner completed more encryptions than the loser.

**3.2.5 Workers' Awareness about Rules.** Finally, we vary the timing of when workers learn how luck plays a role in determining outcomes. In the rules-before condition, we inform workers that effort multipliers or a coin flip can influence the outcome *before* they start the task. In the rules-after condition, we inform workers that multipliers or a coin flip can influence the outcome *after* they complete the task. We randomly assign half of the spectators in the ex-post lucky opportunities and lucky outcomes conditions to the rules-before treatment and half to the rules-after treatment. By construction, we assign all participants in the baseline lucky opportunities condition to the rules-before treatment. Spectators have complete information about when workers learned how we determine the winner.

### 3.3 Comprehension Checks and Belief Elicitation

To ensure that spectators understand the design details, we implement several comprehension questions after they see the initial instructions about the worker task. These questions test spectators' understanding of how luck can affect outcomes and their awareness of when workers learn about the importance of luck. Spectators can only continue once they select the correct answer, and we briefly explain why the answer is correct once they submit it. Therefore, these questions serve as both a comprehension check and as reminders that reinforce the critical aspects of the worker task that are central to our design.

After the 12 redistribution decisions, spectators complete a brief exit survey. The first part of the exit survey consists of three questions. First, we randomly select one of the 12 decisions that the spectators made and present the same information to them. We then ask spectators in the lucky outcomes condition how many encryptions they think workers solved on average. For spectators in the lucky opportunities condition, we randomly draw a multiplier and ask how many encryptions they think workers with that multiplier solved on average. Finally, we asked them how much they would allocate to the winner if they knew they had solved more encryptions.

The second part of the exit survey asks spectators to select their level of agreement with several belief statements in a five-point Likert scale grid. It probes their views on various topics relating to income redistribution and the role of the government. We also embed an attention check in one of the rows that states: "Select disagree if you are reading this."

### 3.4 Recruitment of Workers and Spectators

We recruited 2,416 participants on Amazon Mechanical Turk to participate in the worker task in September 2021. We restricted participation to workers that were U.S.-based, had a 95 percent minimum approval rate, and had at least 500 approved human intelligence tasks (HITs). We excluded 16 participants who completed less than one encryption per minute for a final sample of 2,400 workers. We paid all workers a fixed participation fee of $2 upon task completion. Workers also received an additional payment of up to $5 based on the decision of a randomly chosen spectator approximately six weeks after completing the task.

[Table 1 Here]

We recruited 1,170 panelists from the Federal Reserve Bank of New York's *Survey of Consumer Expectations* (SCE) to participate in the spectator task in October and November 2021. This survey targets a non-convenience, nationally representative panel of U.S. heads of households (Armantier et al., 2017).[13] We conducted two waves of recruitment over two and a half weeks to reach our target sample size. Our recruitment sample includes all recent panelists who have taken part in SCE over the past three years but excludes those who are still actively participating in the SCE core module. We targeted a representative sample of US households based on the same criteria used for entry into the core SCE sample. We invited a total of 3,500 panelists to participate in the experiment over the two waves and had an average response rate of 33.4%.

The experimental interface was mobile-friendly to encourage hard-to-reach demographic groups to participate in our experiment. Given the high-quality nature of the SCE sample, we pre-registered to not impose any exclusion criteria for our spectator sample. Our in-sample data-quality metrics support this decision: The median spectator spent 15 minutes on the survey, 89 percent passed the attention check, and 77 percent passed all four comprehension questions on their first attempt. No spectator failed to answer more than two comprehension questions. We paid all respondents a $5 Amazon gift card for completing the survey.

Table 1 reports descriptive statistics for the spectator sample.[14] The average spectator is 49 years old; 52 percent are female, 62 percent are married, and 14 percent are non-white. More than 62 percent of spectators attained a college degree, 58 percent work full-time, 15 percent work part-time, and 20 percent are retired. About a quarter (24 percent) of spectators have a household income below $40,000 per year and about a third (32 percent) more than $100,000 per year. Our sample includes individuals living in all 50 states plus Washington, DC. About 23 percent of spectators live in the Midwest, 21 percent in the Northeast, 35 percent in the South, and 22 percent in the West. Columns 2–7 show that spectator characteristics are similar across conditions.

## 4 Main Results

The primary outcome we examine is the fraction of earnings, $r_{ip}$, that spectator $i$ redistributes from the winner to the loser in worker pair $p$. We refer to the worker who initially receives the total

[13] The SCE recruitment process works as follows. Prospective panelists receive an initial in-take survey that is distributed via mail to a random sample of US addresses, which contains a $1 bill and offers an additional $5 payment for completing and returning the paper-based survey that collects detailed demographic information and some initial economic expectations. Respondents are informed that they may be eligible to participate in additional surveys and asked to provide their email address if they would like to participate. Stratified random sampling is then used to selectively invite respondents to join the core SCE panel to achieve a nationally representative sample of US household heads in terms of income, gender, age, race/ethnicity, and Census division. Respondents participate in the core SCE survey for up to 12 months and receive a $15 payment for each survey. They can also receive additional payments for completing special surveys, such as the one we fielded for this study. See Section 3 in Armantier et al. (2017) for additional details about the SCE intake and recruitment process.

[14] See Table A2 in the Appendix for summary statistics on our sample of workers. Panel A of Appendix Figure A2 plots the distribution of worker performance.

earnings as the "winner" and the worker who initially receives no earnings as the "loser." When $r_{ip} = 0$, the loser gets none of the total earnings, and the winner retains all the earnings. Both workers receive half of the total earnings when $r_{ip} = 0.5$. The primary outcomes, hypotheses, and empirical specifications described in this section were pre-registered in our pre-analysis plan. We note any deviations in the main text.

### 4.1 Redistribution under Lucky Opportunities and Lucky Outcomes

Spectators redistribute less income when luck is manifested through unequal opportunities. Panel A of Table 2 reports the average level of redistribution across our luck treatments. Under lucky outcomes, spectators redistributed 27.6 percent of earnings from the winner to the loser on average. In contrast, spectators redistributed only 23.4 percent of earnings on average under lucky opportunities. Thus, spectators redistributed 4.2 percentage points less of total income when luck was experienced through unequal opportunities ($p < 0.01$, column 3). This difference equates to a 15.3 percent decrease in the final earnings for the losing worker. If we restrict our analysis to observations in which spectators face uncertainty about whether the inequality is merit-based, we find that spectators redistributed 5 percentage points ($p < 0.01$) less of total income under lucky opportunities, implying 17.3 percent lower earnings for the losing worker.[15]

Next, we examine how redistribution responds to the importance of luck in determining outcomes. Figure 1 plots mean redistribution in each luck treatment against the probability that the best performer lost: $1 - \pi$. Intuitively, $1 - \pi$ can be thought of as the probability that initial earnings inequality was due to luck or, equivalently, the probability that earnings inequality is not merit-based. We depict $1 - \pi$ on the x-axis, rather than $\pi$, because this makes it easier to discuss the effect of increasing the impact of luck, which facilitates the exposition of our results. The intercept in Figure 1 shows mean redistribution across worker matches when the best performer was always the winner— because workers drew equal multipliers or because there was zero probability of earnings being determined by a coin flip. The slope shows how redistribution changes when there is an increase in the probability that initial earnings inequality was due to luck.

[Figure 1 Here]

We first examine worker pairs where performance alone determined the winner, eliminating any uncertainty about whether initial earnings inequality stemmed from worker performance or luck. In these merit-based outcomes, spectators behave similarly across luck environments.[16]

[15] We refer to situations in which $\pi = 1$ as situations without uncertainty because the information available to spectators in those situations allows them to conclude without a doubt that the winner of a worker pair is the best performer of the pair. For example, this is the case when the coin flip probability is zero in the lucky outcomes condition or when both workers have the same multiplier in the lucky opportunities condition. Upon excluding these observations, average redistribution equals 28.9 percent in the lucky outcomes environment versus 23.9 percent in the lucky opportunities environment.

[16] Notably, redistribution is slightly *higher* under lucky opportunities in these merit-based outcomes, contrary to the overall redistribution gap documented in Panel A of Table 2. This pattern may partly reflect spectator confusion

We next examine how redistribution responds to an increase in the probability that the best performer lost. Under lucky outcomes, spectators' redistribution decisions initially react strongly to this probability. As the probability that the best performer lost increases from 0 to about 25–30 percent, the share of total income that spectators redistribute increases from 13.1 to 34.5 percent. Beyond that point, redistribution does not respond to further increases in the probability that the best performer lost. In contrast, spectators' reactions are more muted when there are lucky opportunities. As the probability that the best performer lost increases from 0 to the range 25–30 percent, average redistribution increases from 17.9 to 22.8 percent. However, spectators continue to respond to the increased role that luck plays in determining outcomes beyond that point, increasing the amount they redistribute in an approximately linear manner.

To summarize the relationship between average redistribution and the likelihood that the best performer in match $p$ lost in each luck environment, we estimate linear models that relate the share of earnings redistributed to $1 - \pi_{ip}$:

$$r_{ip} = \alpha + \beta(1 - \pi_{ip}) + \varepsilon_{ip}, \tag{7}$$

where $\varepsilon_{ip}$ is an error term. The main parameter of interest is $\beta = \partial \mathbb{E}(r_{ip})/\partial \mathbb{E}(1 - \pi_{ip})$, which measures the elasticity of redistribution with respect to $1 - \pi_{ip}$. The exogenous within-subject variation in $\pi_{ip}$ allows us to identify $\beta$.

Spectators respond less to changes in the probability that the outcome was luck-based when luck stems from unequal opportunities instead of coin flips. Panel B of Table 2 presents estimates of $\beta$ across the different luck environments. A ten percentage point increase in $1 - \pi_{ip}$ leads to a $\hat{\beta} = 3.7$ percentage point increase in the share of earnings redistributed in the lucky outcomes condition (column 1). In contrast, redistribution is less elastic to changes in $1 - \pi_{ip}$ in the lucky opportunities condition: A ten percentage point increase in $1 - \pi_{ip}$ leads to a $\hat{\beta} = 2.0$ percentage point increase in redistribution when luck emerges through unequal opportunities (column 2). The difference between estimated elasticities is statistically significant ($p < 0.01$, column 3).[17]

This differential sensitivity to changes in $1 - \pi_{ip}$ leads to a substantial redistribution gap across our experimental treatments. Column (3) of panel C in Table 2 presents the average difference in

in cases where the winner had the *lower* multiplier. While winners had higher multipliers in 11 of 12 decisions, for the merit-based outcome case ($1 - \pi = 0$), the winning worker had either a lower or equal multiplier. When the winning worker had the lower multiplier, some spectators may not have noticed this difference from their other decisions. Indeed, when the multipliers were exactly equal, the average redistribution was 10.7 percent, compared to 19.2 percent when the winner had a lower multiplier. Although the confidence intervals are relatively large when we analyze these sub-samples separately, this suggests that some spectators may have erroneously attributed the lower multiplier to the loser for the pure merit decision.

[17] These results remain robust when we exclude the two extreme cases $1 - \pi_{ip} = 0$ and $1 - \pi_{ip} = 0.50$. When excluding all observations that correspond to these extreme cases from the analysis, we find that a ten percentage point increase in $1 - \pi_{ip}$ leads to a 3.5 percentage point increase in redistribution under lucky outcomes and a 1.8 percentage point increase under lucky opportunities, maintaining the 1.7 percentage point gap in responsiveness between the two environments.

redistribution for each $\pi$ bin, along with the estimated standard error.[18] We observe a significant difference in redistribution for six out of the ten $\pi$ bins in which spectators face inferential challenges when trying to determine whether the inequality was merit-based ($p < 0.05$).[19]

[Table 2 Here]

### 4.2 Margins of Redistribution and Heterogeneity

Differences in redistribution across luck environments can be decomposed into differences in the extensive margin (whether anything is redistributed) and the intensive margin (how much is redistributed).[20] In Appendix B.4, we investigate the role of the extensive and intensive margins in accounting for the redistribution gap across luck environments. On the extensive margin, we find that a higher fraction of spectators do not redistribute any earnings under lucky opportunities than under lucky outcomes ($p < 0.01$). On the intensive margin, we still find that redistribution is less responsive to the likelihood that the best performer lost under lucky opportunities than lucky outcomes ($p < 0.01$). Thus, both the extensive and the intensive margin contribute to the redistribution gap.

We also find that redistribution behavior varies by respondents' gender, income, and political views. Appendix Table A3 presents the results from a heterogeneity analysis. Females tend to favor redistribution more than males across both luck environments. Respondents in households with annual incomes below $100,000 redistribute more on average across both luck environments. This finding is notable in light of recent cross-country evidence by Maréchal et al. (2025) suggesting that actual redistribution policies are more strongly correlated with the preferences of lower socioeconomic groups. Finally, Republicans tend to support redistribution less than others. These patterns are consistent with previous evidence and extend to the domain of unequal opportunities (Alesina and Giuliano, 2011; Ashok, Kuziemko, and Washington, 2015; Alesina, Stantcheva, and Teso, 2018; Almås, Cappelen, and Tungodden, 2020; Cohn et al., 2023).

[18] Formally, in Panel C of Table 2, we estimate regressions of the form:

$$r_{ib} = \sum_{b=1}^{12} \gamma_b (1 - \pi_b) + \varepsilon_{ib}, \tag{8}$$

where $1 - \pi_b$ is an indicator that equals one if $1 - \pi_{ip}$ is in bin $b$. We estimate equation (8) separately for each treatment and interact the bins with treatment dummies to assess formally whether mean redistribution is the same across luck treatments at a given $\pi$ bin. We cluster standard errors at the spectator level in all specifications.

[19] When making decisions for two of the twelve $\pi$ bins (i.e., in the two extreme cases: $1-\pi_{ip} = 0$ and $1-\pi_{ip} = 0.50$), spectators trying to assess whether merit could have determined the winner do not face the same inferential challenges as in the other ten decisions. This is because $1 - \pi_{ip} = 0$ corresponds to cases in which the probability of a coin flip is zero (in lucky outcomes) or the winner does not have any relative advantage (in lucky opportunities) so that the winner is for sure the better performer. Similarly, $1-\pi_{ip} = 0.5$ corresponds to cases in which winning is independent of performance because the coin flip probability is 100% (in lucky outcomes) or the winning worker's relative advantage was too large to be overcome (in lucky opportunities).

[20] This analysis of the intensive and extension margins of redistribution was not included in our preregistration and should therefore be considered exploratory in nature.

### 4.3 The Impact of Observable Features on Redistribution

To understand why spectators are less responsive to changes in the probability of the best performer losing under lucky opportunities compared to lucky outcomes, we conduct a two-step analysis. First, we examine how spectators incorporate observable features of the luck environment into their redistribution decisions (Section 4.3.1). Second, we analyze the relationship between these features and the probability that the best performer lost (Section 4.3.2).

**4.3.1 Spectator Responses to Unequal Opportunities.** We begin by analyzing how spectators incorporate unequal opportunities into their redistribution decisions. We create two intuitive measures of unequal opportunities: (1) the arithmetic difference between the productivity multipliers of the winner and loser ($m_1 - m_2$), and (2) the ratio of these multipliers ($m_1/m_2$). We then evaluate which measure better predicts spectators' redistribution choices by regressing the fraction of earnings redistributed on each measure (Table 3). We find that when either the multiplier difference or ratio is included in the regression by itself, each has a significant impact on redistribution (columns 1 and 2). However, when both measures are included simultaneously in the regression model, only the multiplier difference continues to have a significant effect (column 3).[21]

These results suggest that spectators primarily rely on the multiplier difference (and not the ratio) when making redistribution decisions in the lucky opportunities environment. For example, spectators treat multipliers $(m_1, m_2) = (2, 1)$ the same as $(m_1, m_2) = (3, 2)$ when making redistribution decisions, when in fact it is more likely that outcomes were luck-based in the former; specifically, $1 - \pi(2/1) = 0.46$ compared to $1 - \pi(3/2) = 0.38$. A worker competing against another with a multiplier of one would only require a multiplier of 1.5 to have an equivalent advantage to a worker with a multiplier of three competing against one with a multiplier of two.

[Table 3 Here]

Figure 2 provides visual evidence of how spectators use the features from the environment in each luck condition to inform their redistribution decisions. The figure presents a series of binned scatterplots plotting average redistribution ($y$-axis) against the information observed by spectators ($x$-axis). Under lucky outcomes (Panel A), the relationship between redistribution and the coin-flip probability is piecewise linear. For low coin-flip probabilities, there is a positive and linear relationship, but for sufficiently large coin-flip probabilities, redistribution is unresponsive to further changes in the coin-flip probability. Under lucky opportunities (Panel B), the relationship between redistribution and the multiplier difference is linear across the entire range of multiplier differences.[22] On average, a one-unit increase in the difference between the high and low performers'

[21] The coefficient on the multiplier difference in column 3 is not statistically different from the corresponding estimate in column 1 ($p = 0.42$). However, we can reject that the coefficient on the multiplier ratio is the same in columns 2 and 3 ($p < 0.01$).

[22] In Appendix Table A10, we regress the fraction of earnings redistributed on higher-order polynomials of the multiplier difference and find that none of these terms have a significant effect on redistribution. This provides further evidence suggesting that spectators use the *linear* multiplier difference to inform their redistribution decisions.

multipliers leads to a four percentage point increase in the share of earnings redistributed.

[Figure 2 Here]

**4.3.2 Features of the Environment and Source of Inequality.** Next, we study the empirical relationship between the observable features of each luck environment and the probability that the inequality between workers is the result of luck. Under lucky outcomes, there is a linear relationship between the coin-flip probability and the likelihood that the best performer lost (Figure 3, Panel A). This relationship starts at the origin, as the best performer never loses when the coin-flip probability is zero, and increases linearly with the coin-flip probability.

In contrast, under lucky opportunities, there is a concave relationship between the multiplier difference and the likelihood that the best performer lost (Figure 3, Panel B). When multipliers are equal, the best performer never loses. However, even a small increase in the multiplier difference leads to a steep rise in the probability of the best performer losing. For example, increasing the multiplier difference from 0 to 0.5 increases the likelihood of the best performer losing from 0 to about 20 percent. As the multiplier difference becomes sufficiently large, the best performer lost whenever they did not receive the high multiplier.

[Figure 3 Here]

Taken together, the combination of how redistribution responds to the environmental features and how these features relate to the probability that the best performer lost explains why spectators underreact under unequal opportunities. Under lucky outcomes, the coin-flip probability and the likelihood that the best performer lost are linearly related (Figure 3, Panel A). As a result, the relationship between redistribution and the coin-flip probability (Figure 2, Panel A) exactly mirrors the relationship between redistribution and the probability that the best performer lost (Figure 1, black line).

In the lucky opportunities condition, spectators' redistribution decisions respond linearly to the multiplier difference across its entire range (Figure 2, Panel B). However, this linear response fails to account for the concave relationship between the multiplier difference and the probability that the best performer lost (Figure 3, Panel B). As a result, spectators appear to "underreact" (redistribute less than expected based on the actual likelihood that the best performer lost) for relatively small values of $(1-\pi)$ and "overreact" for larger values (Figure 1, grey line).[23] In Section 5, we investigate whether this linear response to multiplier differences persists when we also provide direct information about $\pi$.

[23] In Online Appendix Section B.3, we show that the same relationship prevails even if lucky opportunities are not multiplicative but additive. More precisely, we consider the possibility that workers receive unequal, additive boosters, capturing the idea that a relative advantage can resemble a head start sometimes. In such an environment, the probability that the inequality is merit-based is still convex in the relative advantage (the difference in boosters), implying that $(1-\pi)$ is concave in the booster difference – as is the case for our design using productivity multipliers.

## 5 The Role of Information and Effort Responses

We explore two broad categories of mechanisms that may drive the patterns of redistribution that we observe across luck environments. First, we investigate the role of inaccurate beliefs about the probability the best performer lost through an information intervention. Second, we examine the role of actual or perceived differences in worker performance across luck environments in explaining our results.

### 5.1 Information Intervention

A key question is why spectators are less responsive to changes in the likelihood that the winner was luck-based in lucky opportunities versus lucky outcomes, which leads to substantial gaps in average redistribution across these environments. To examine whether this is due to spectators facing inferential challenges when assessing the impact of unequal opportunities on earnings inequality, we implement an information treatment. In this treatment, in addition to the information spectators receive in the baseline condition, we provide spectators with the information they would otherwise need to infer. Specifically, for each of their 12 redistribution decisions, we tell spectators the likelihood that the best performer won, $\pi$.

We begin by examining how the information intervention affects the relationship between redistribution and the likelihood that the outcome is luck-based. If inferential challenges are the primary driver of the linear redistribution pattern, then we expect the treatment to significantly alter this relationship, as spectators no longer need to rely on heuristics or potentially inaccurate inference. In contrast, if the linear pattern persists, this would suggest that other factors shape redistributive behavior. Then, we analyze how the information intervention affects spectators' reliance on observable features of the luck environment.

[Figure 4 Here]

Providing readily accessible information about $\pi$ made redistribution more responsive to the likelihood that the best performer lost in both luck environments. Figure 4 plots mean redistribution in the baseline and information interventions. In the baseline lucky outcomes condition, a ten percentage point increase in the likelihood that the best performer lost raised redistribution by 3.7 percentage points. In the information-treated group, this effect increases by 1.5 percentage points ($p < 0.05$, Appendix Table A4, Panel B) to 5.2 percentage points in total. Similarly, in the baseline lucky opportunities condition, a ten percentage point increase in the likelihood that the best performer lost raised redistribution by 2.0 percentage points. In the information-treated group, this effect increases by 1.7 percentage points ($p < 0.01$) to 3.7 percentage points in total, implying a 85 percent relative increase in the elasticity of redistribution. In other words, when information about the probability that earnings inequality was due to luck becomes readily accessible, spectators

become more sensitive to this probability and adjust their redistribution decisions accordingly.[24]

As a consequence of this increased sensitivity, the information treatment reduced the redistribution gap across luck environments. Figure 5 plots the redistribution gap between the lucky outcomes and lucky opportunities environment in the baseline treatments with no information about $\pi$ (black line) and readily accessible information about $\pi$ (grey line). The redistribution gap diminishes substantially when information about $\pi$ is provided, with the most pronounced declines observed over the range $1-\pi \in [20, 40)$. A joint F-test of the null hypothesis that the change in the redistribution gap between lucky opportunities and lucky outcomes across all $\pi$ bins equals zero is rejected at the 1% significance level (F=2.19, p=0.01). Notably, the gap is small and statistically indistinguishable from zero for ten of the twelve $\pi$ bins, and is statistically significant at $p < 0.05$ for only two $\pi$ bins (Appendix Table A5, column 2). This finding is consistent with the interpretation that spectators face inferential challenges when making redistribution decisions.

[Figure 5 Here]

Next, we investigate how the information intervention affected the spectators' reliance on the other observable features of the two luck environments.[25] Figure 6 shows a series of binned scatterplots analogous to those in Figure 2. Circle markers reproduce the relationship between mean redistribution and the observable features in the baseline luck environments. Triangle markers show this same relationship but under information treatments. Finally, square markers show the relationship between mean redistribution and the observable information, after removing the influence of $\pi$ on redistribution.[26] Solid lines plot the predicted values from local polynomial estimates.

[Figure 6 Here]

The information treatment altered how spectators incorporated observable features of the luck environment into their redistribution decisions. Figure 6 illustrates how the relationship between redistribution and coin-flip probabilities/multiplier differences changed once information on the likelihood that inequality is merit-based became readily available to spectators. In the lucky outcomes environment, the information treatment had two effects: it decreased mean redistribution for most coin-flip probabilities, and it increased the responsiveness of redistribution for relatively high coin-flip probabilities. In the lucky opportunities environment, the information treatment reduced

[24] The effects of providing information about $\pi$ are robust to again excluding the two extreme situations $1-\pi_{ip} = 0$ and $1-\pi_{ip} = 0.50$ where spectators face less inferential challenges. Excluding these observations from the analysis of the information treatments, we find that a ten percentage point increase in the likelihood that the best performer lost raises redistribution by 4.8 percentage points under lucky outcomes (versus 3.5 percentage points in the baseline condition) and by 3.6 percentage points under lucky opportunities (versus 1.8 percentage points in the baseline condition).

[25] We did not pre-register this analysis; however, our linearity results from Section 4.3.2 invite further exploration.

[26] More precisely, we plot the residual of an individual's redistribution after subtracting from it the component explained by the likelihood that the best performer lost. To construct these residuals, we first regress the share of earnings redistributed on the likelihood that the best performer lost (as in equation (7)) and estimate the residuals from this regression, adding back the unconditional sample mean to facilitate the interpretation of units.

mean redistribution for relatively low multiplier differences. Moreover, it made redistribution more responsive to changes in the multiplier difference for relatively small multipliers, particularly in the range where a small increase in relative multipliers has a substantial impact on worker outcomes.

The relationship between mean redistribution and the observable features of the luck environment appears to be mediated by the likelihood that the outcome was merit-based. Controlling for the likelihood that the outcome is merit-based attenuates the relationship between redistribution and the features of the environment observed by spectators (Figure 6, gray lines). Once this likelihood is controlled for, these features have no independent impact on redistribution.[27] This finding further supports the hypothesis that spectators primarily use observable features of the environment as proxies to infer the likelihood of merit-based outcomes.

### 5.2 Timing of Luck and Worker Performance

An important difference between our luck environments is that lucky outcomes occur after completing the task, while unequal opportunities are known before. This aspect of the lucky opportunities condition reflects how luck typically arises in many real-life situations. This difference in the timing of luck could drive the differences in the redistribution decisions that we observe if spectators have different expectations about how workers respond to getting a high or low multiplier.[28]

To examine whether the timing of lucky opportunities affects redistribution, we implemented an additional between-subjects treatment that aligned the timing of when workers learn about their luck across environments. In the *ex-post lucky opportunities* condition, workers learn their multipliers only after they complete the task. This is in contrast to our baseline lucky opportunities condition, in which we inform workers of their multipliers before they begin working on the encryption task. However, this timing is in line with when workers learn about the probability that a coin flip determines their earnings in the lucky outcomes environment, namely after completing the task. To ensure this intervention is known to spectators, we provide them a visual timeline of the worker task sequence and require that they pass a comprehension check about whether multipliers are revealed to workers before or after the task.[29]

We find that whether opportunities are encountered before or after working has no economically

[27] Appendix Table A11 provides the corresponding regression estimates and confirms that the multiplier difference has no significant effect on redistribution when controlling for $\pi$ in the information intervention.

[28] What matters for redistribution behavior are spectators' beliefs about worker performance, which we control for using our additional treatments. Empirically, we do not find any evidence that worker performance responds to receiving a high or low multiplier (see Table A7 in the Appendix). We also observe no differences in the overall performance distribution across luck environments (see Appendix Figure A3). A Kolmogorov–Smirnov test for equality of distribution cannot reject the hypothesis that the distribution of worker performance in the lucky outcomes and lucky opportunities environments are equal ($p = 0.909$). Notably, other work has found performance responses to differences in opportunities (Andre, 2024). This difference likely arises because our environment is a winner-takes-all tournament with a fixed working period, while Andre (2024) considers differential piece-rate wages and allows workers to choose how long they work. Indeed, DellaVigna et al. (2022) find that higher incentives lead to higher output when workers can choose how long they work for but have no effect when there is a fixed working period.

[29] The equivalent visual timelines and comprehension checks are provided to all spectator conditions, and not only those spectators in ex-post lucky opportunities. See Supplementary Materials, Figures 5 and 18.

or statistically significant impact spectators' redistribution decisions. In Appendix Table A8, we re-estimate the primary specifications in Table 2 but compare redistribution between the baseline and ex-post lucky opportunities. We find no significant differences in the average redistribution: The average amount of income redistributed was 23.4 percent in baseline lucky opportunities versus 24.4 percent in ex-post lucky opportunities ($p = 0.57$). We also find no significant differences in the elasticity of redistribution to changes in luck ($p = 0.89$). Figure 7 plots our estimates of the average redistribution for both lucky opportunities conditions across each $\pi$ bin. Across the entire range of $\pi$ bins, we find no differences in the level of redistribution.[30]

[Figure 7 Here]

To further assess whether spectators expected differential effort from workers who received a low versus high multiplier, we elicited their stated beliefs about worker performance across the multiplier distribution. To do this we randomly selected a multiplier for each spectator in the baseline lucky opportunities condition and elicited their beliefs about the average number of encryptions completed by workers who received that multiplier. This information enables us to recover the perceived elasticity of worker performance with respect to the productivity multiplier, which we estimate by regressing spectator beliefs about worker performance on the randomly selected multiplier (Appendix Table A9). We find no evidence that spectators expect a significant worker performance response from receiving a high or low multiplier.[31]

Finally, it is possible that spectators believed that merely knowing there would be coin flips or multipliers causes workers to put in differential effort across our luck environments, even if the exact coin flip probability or productivity multiplier is revealed to workers only after completing the task. To fully eliminate any scope for differential beliefs about workers' performance responses, we implemented additional subtreatments in which workers did not receive any information about how the winner would be determined *prior* to working on the task. Crucially, spectators in this "rules-after" scenario in both the lucky outcomes and ex-post lucky opportunities conditions knew that workers faced identical instructions before completing the task. We analyze the data from these subtreatments in Online Appendix Section B.5. In Appendix Figure B3, we show that whether workers learn about the possibility of a coin flip or multipliers affecting the outcome before or after working on the task has no impact on spectators' support for redistribution. In addition, Appendix Figure B4 and the estimates in Appendix Table B3 show that there is still a redistribution gap as well as a difference in the elasticity of redistribution between lucky outcomes and lucky opportunities

[30] Accordingly, there is still a redistribution gap as well as a difference in the elasticity of redistribution between the ex-post lucky opportunities condition and the lucky outcomes condition. For example, if the likelihood that inequality is due to luck increases by ten percentage points, redistribution increases by 1.6 percentage points ($p < 0.01$) less under ex-post lucky opportunities (versus 1.7 percentage points in our main comparison in Section 4.1).

[31] Spectators could also expect that the difference in performance between the winner and loser is not identical across our luck environments. In Appendix B.6, we show that the effort gap between winners and losers is very similar in the two luck conditions across the entire range of $\pi$ bins. Moreover, we provide evidence that potential differences in the winner-loser effort gap cannot explain the redistribution gap between the two luck environments.

in the rules-after treatment, that is, when workers had identical information before completing the task.

## 6 Predicting Political and Social Views

We document that individuals redistribute less when luck arises through unequal opportunities rather than exogenous sources. But which environment is more predictive of the spectator's political and social views? Our data allows us to examine the external relevance of the experimental measure of redistribution attitudes based on our lucky opportunities environment and compare it to measures based on the lucky outcomes environment used in prior work. To assess this, we estimate the correlation between the average earnings redistributed by spectators in each environment and 14 self-reported political and social views.[32]

Redistribution behavior under lucky opportunities correlates more with real-world social and political attitudes than the lucky outcomes environment. We find that mean redistribution in the lucky opportunities condition is more predictive than in lucky outcomes for 13 out of the 14 social and political attitudes we measure (Appendix Table A6). To create summary measures of political attitudes, we calculate a $z$-score that averages the 14 self-reported individual attitudes and an index that equals the first component of a principal component analysis. The summary indices reveal that behavior in the lucky opportunities condition is 45 to 62 percent more predictive of attitudes than in the lucky outcomes condition. While the overall explanatory power is modest, with R-squared values ranging from 1.4 to 4.4 percent, the lucky opportunities environment consistently explains nearly three times as much variation in political attitudes as the lucky outcomes environment. This finding suggests that redistribution decisions made in the lab better reflect real-world social and political views when luck arises through unequal opportunities rather than unequal outcomes.

## 7 Discussion

Meritocratic fairness ideals contend that individuals are willing to tolerate inequalities due to differences in effort but oppose those arising from chance. But when people face unequal opportunities, this distinction is obfuscated by the fact that luck and effort are intertwined. As a result, individuals may find it difficult to assess the source of inequality and may misattribute success to effort rather than luck. This paper shows that impartial third parties are more tolerant of inequality when luck arises through unequal opportunities, and there is uncertainty about the role of luck in determining outcomes.

We provide evidence that a key mechanism driving these results is the differential ability to understand the impact of luck across environments, as spectators face greater inferential challenges

[32]We collected participants' social and political views to conduct validation of our experimental measures and heterogeneity analysis but did not pre-register specific hypotheses. As a result, we consider the analysis in this section as exploratory.

in mapping features of the luck environment to their consequences for outcomes when luck manifests through unequal opportunities.

Our results offer a potential rationale for the apparent disconnect between the previous experimental literature and observed patterns of U.S. inequality. Research that generates inequality through exogenous variation in outcomes has found that most Americans equalize incomes when income differences are due to luck (Almås, Cappelen, and Tungodden, 2020). However, support for redistribution in the U.S. remained stagnant over a period when differences in opportunities became increasingly important (Chetty et al., 2014; Ashok, Kuziemko, and Washington, 2015). Consistent with these trends, we show that redistribution is less sensitive to changes in luck when luck arises through unequal opportunities. Similarly, the U.S. remains the most unequal country in the OECD while ranking poorly on equality of opportunity (Corak, 2013; Mitnik, Helsø, and Bryant, 2020). Consistent with these cross-country comparisons, we show that Americans tolerate more inequality when it arises due to differential opportunities.

We conclude by discussing two implications of our results for models that seek to understand and predict attitudes toward redistribution. First, we document that in the absence of readily available information about the likelihood that outcomes are merit-based, spectators rely on simple heuristics when factoring the impact of luck into their redistribution decisions. As a result, people fail to appreciate how small differences in initial circumstances can greatly impact outcomes. Providing information about the importance of luck reduces this reliance on heuristics. Models that seek to accommodate cognitive errors hold some promise for predicting and explaining how beliefs shape redistribution attitudes.

Second, our results show that readily available information about the likelihood that outcomes are merit-based greatly impacts people's redistribution decisions. This suggests that the information individuals frequently encounter might disproportionately impact their views on inequality and redistribution. For example, popular media coverage (e.g., rags-to-riches stories) may lead individuals to have a greater tolerance for inequality of opportunities, making them less willing to correct this source of unfairness through redistribution. Exploring how salient information shapes individuals' tolerance for inequality is a promising avenue for future research.

Taken together, our results highlight that redistributive behavior is not invariant to how luck is implemented in the lab. The lucky opportunities environment has several important features that affect redistribution, which the simpler lucky outcomes paradigm cannot capture. We provide a portable, tractable, and rich environment to study income redistribution when there are unequal opportunities that can inform the development of inequality models and the design of optimal redistribution policies.

## Data Availability

Data and code for replicating the tables and figures in this article can be found in Preuss et al. (2025) in the Harvard Dataverse, URL: `https://doi.org/10.7910/DVN/H7QOXG`.

## References


Alesina, Alberto and Paola Giuliano. 2011. "Preferences for Redistribution." In *Handbook of Social Economics*, vol. 1, edited by Jess Benhabib, Alberto Bisin, and Matthew O. Jackson, chap. 4. North-Holland, 93–131.

Alesina, Alberto, Stefanie Stantcheva, and Edoardo Teso. 2018. "Intergenerational Mobility and Preferences for Redistribution." *American Economic Review* 108 (2):521–54.

Almås, Ingvild, Alexander W. Cappelen, and Bertil Tungodden. 2020. "Cutthroat Capitalism versus Cuddly Socialism: Are Americans More Meritocratic and Efficiency-Seeking than Scandinavians?" *Journal of Political Economy* 128 (5):1753–88.

Andre, Peter. 2024. "Shallow Meritocracy." *Review of Economic Studies* 00 (0):1–36.

Armantier, Olivier, Giorgio Topa, Wilbert Van der Klaauw, and Basit Zafar. 2017. "An Overview of the Survey of Consumer Expectations." *Economic Policy Review* 23 (2):51–72.

Ashok, Vivekian, Ilyana Kuziemko, and Ebonya Washington. 2015. "Support for Redistribution in an Age of Rising Inequality: New Stylized Facts and Some Tentative Explanations." *Brookings Papers on Economic Activity* 46 (1):367–405.

Bagnoli, Mark and Ted Bergstrom. 2006. "Log-concave Probability and its Applications." In *Rationality and Equilibrium*, *Studies in Economic Theory*, vol. 26, edited by Charalambos D. Aliprantis, Rosa L. Matzkin, Daniel L. McFadden, James C. Moore, and Nicholas C. Yannelis. Springer, 217–41.

Benjamin, Daniel J. 2019. "Errors in Probabilistic Reasoning and Judgment Biases." In *Handbook of Behavioral Economics: Applications and Foundations 1*, vol. 2, edited by B. Douglas Bernheim, Stefano DellaVigna, and David Laibson, chap. 2. North-Holland, 69–186.

Benndorf, Volker, Holger A. Rau, and Christian Sölch. 2019. "Minimizing Learning in Repeated Real-effort Tasks." *Journal of Behavioral and Experimental Finance* 22:239–48.

Bhattacharya, Puja and Johanna Mollerstrom. 2022. "Lucky to Work." GMU Working Paper in Economics. No. 22-46.

Bozio, Antoine, Bertrand Garbinti, Jonathan Goupille-Lebret, Malka Guillot, and Thomas Piketty. 2024. "Predistribution versus Redistribution: Evidence from France and the United States." *American Economic Journal: Applied Economics* 16 (2):31–65.

Cappelen, Alexander W., Astri Drange Hole, Erik Ø. Sørensen, and Bertil Tungodden. 2007. "The Pluralism of Fairness ideals: An Experimental Approach." *American Economic Review* 97 (3):818–27.

Cappelen, Alexander W, James Konow, Erik Ø. Sørensen, and Bertil Tungodden. 2013. "Just Luck: An Experimental Study of Risk-taking and Fairness." *American Economic Review* 103 (4):1398–413.

Cappelen, Alexander W., Karl Ove Moene, Siv-Elisabeth Skjelbred, and Bertil Tungodden. 2022a. "The Merit Primacy Effect." *The Economic Journal* 133 (651):951–70.

Cappelen, Alexander W, Johanna Mollerstrom, Bjørn-Atle Reme, and Bertil Tungodden. 2022b. "A Meritocratic Origin of Egalitarian Behaviour." *The Economic Journal* 132 (646):2101–17.

Cappelen, Alexander W, Erik Ø. Sørensen, and Bertil Tungodden. 2010. "Responsibility for What? Fairness and Individual Responsibility." *European Economic Review* 54 (3):429–41.

Charness, Gary and Matthew Rabin. 2002. "Understanding Social Preferences with Simple Tests." *The Quarterly Journal of Economics* 117 (3):817–69.

Chen, Daniel L., Martin Schonger, and Chris Wickens. 2016. "oTree–An Open-source Platform for Laboratory, Online, and Field Experiments." *Journal of Behavioral and Experimental Finance* 9:88–97.

Chetty, Raj, Nathaniel Hendren, Patrick Kline, Emmanuel Saez, and Nicholas Turner. 2014. "Is the United States Still a Land of Opportunity? Recent Trends in Intergenerational Mobility." *American Economic Review* 104 (5):141–47.

Cohn, Alain, Lasse J. Jessen, Marko Klašnja, and Paul Smeets. 2023. "Wealthy Americans and Redistribution: The Role of Fairness Preferences." *Journal of Public Economics* 225:104977.

Conlon, John J. 2025. "Attention, Information, and Persuasion." Working Paper.

Corak, Miles. 2013. "Income Inequality, Equality of Opportunity, and Intergenerational Mobility." *Journal of Economic Perspectives* 27 (3):79–102.

Cruces, Guillermo, Ricardo Perez-Truglia, and Martin Tetaz. 2013. "Biased Perceptions of Income Distribution and Preferences for Redistribution: Evidence from a Survey Experiment." *Journal of Public Economics* 98:100–12.

DellaVigna, Stefano, John A. List, Ulrike Malmendier, and Gautam Rao. 2022. "Estimating Social Preferences and Gift Exchange at Work." *American Economic Review* 112 (3):1038–74.

Dong, Lu, Lingbo Huang, and Jaimie W. Lien. 2024. "'They Never had a Chance': Unequal Opportunities and Fair Redistributions." *The Economic Journal* 00 (0):ueae099.

Durante, Ruben, Louis Putterman, and Joël van der Weele. 2014. "Preferences for Redistribution and Perception of Fairness: An Experimental Study." *Journal of the European Economic Association* 12 (4):1059–86.

Enke, Benjamin. 2020. "What You See is All There Is." *The Quarterly Journal of Economics* 135 (3):1363–98.

Erkal, Nisvan, Lata Gangadharan, and Nikos Nikiforakis. 2011. "Relative Earnings and Giving in a Real-Effort Experiment." *American Economic Review* 101 (7):3330–48.

Fehr, Dietmar, Daniel Müller, and Marcel Preuss. 2024. "Social Mobility Perceptions and Inequality Acceptance." *Journal of Economic Behavior & Organization* 221:366–84.

Frank, Robert H. 2016. *Success and Luck: Good Fortune and the Myth of Meritocracy.* Princeton University Press.

Graeber, Thomas. 2023. "Inattentive Inference." *Journal of the European Economic Association* 21 (2):560–92.

Konow, James. 2000. "Fair Shares: Accountability and Cognitive Dissonance in Allocation Decisions." *American Economic Review* 90 (4):1072–92.

Kuziemko, Ilyana, Michael I. Norton, Emmanuel Saez, and Stefanie Stantcheva. 2015. "How Elastic Are Preferences for Redistribution? Evidence from Randomized Survey Experiments." *American Economic Review* 105 (4):1478–508.

Larrick, Richard P. and Jack B. Soll. 2008. "The MPG Illusion." *Science* 320 (5883):1593–94.

Levy, Matthew and Joshua Tasoff. 2016. "Exponential-growth Bias and Lifecycle Consumption." *Journal of the European Economic Association* 14 (3):545–83.

Maréchal, Michel André, Alain Cohn, Jeffrey Yusof, and Raymond Fisman. 2025. "Whose Preferences Matter for Redistribution? Cross-Country Evidence." *Journal of Political Economy Microeconomics* 3 (1):000–000.

Martínez-Marquina, Alejandro, Muriel Niederle, and Emanuel Vespa. 2019. "Failures in Contingent Reasoning: The Role of Uncertainty." *American Economic Review* 109 (10):3437–74.

Mirrlees, J. A. 1971. "An Exploration in the Theory of Optimum Income Taxation." *The Review of Economic Studies* 38 (2):175–208.

Mitnik, Pablo, Anne-Line Helsø, and Victoria L Bryant. 2020. "Inequality of Opportunity for Income in Denmark and the United States: A Comparison Based on Administrative Data." Working Paper No. 27835, National Bureau of Economic Research.

Mollerstrom, Johanna, Bjørn-Atle Reme, and Erik Ø. Sørensen. 2015. "Luck, Choice and Responsibility – An Experimental Study of Fairness Views." *Journal of Public Economics* 131:33–40.

Preuss, Marcel, Germán Reyes, Jason Somerville, and Joy Wu. 2025. "Replication Data for: Inequality of Opportunity and Income Redistribution." https://doi.org/10.7910/DVN/H7QOXG, Harvard Dataverse.

Rees-Jones, Alex and Dmitry Taubinsky. 2020. "Measuring 'Schmeduling'." *The Review of Economic Studies* 87 (5):2399–2438.

Table 1: Average spectator characteristics by treatment condition

| | | Baseline condition | | | Information Treatment | | |
|---|---|---|---|---|---|---|---|
| | All | Lucky Outcomes | Lucky Opportunities | Ex-Post Lucky Opportunities | Lucky Outcomes | Lucky Opportunities | Ex-Post Lucky Opportunities |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
| **Panel A. Demographic characteristics and race** | | | | | | | |
| Age | 49.09 | 49.50 | 50.06 | 47.28 | 48.68 | 49.63 | 49.37 |
| Male | 0.48 | 0.50 | 0.47 | 0.48 | 0.48 | 0.52 | 0.45 |
| Married | 0.62 | 0.61 | 0.63 | 0.68 | 0.57 | 0.66 | 0.60 |
| Nr. of children under 18 | 0.60 | 0.51 | 0.57 | 0.61 | 0.53 | 0.74 | 0.62 |
| White | 0.86 | 0.82 | 0.88 | 0.84 | 0.84 | 0.88 | 0.90 |
| Black | 0.08 | 0.12 | 0.06 | 0.07 | 0.07 | 0.07 | 0.07 |
| Hispanic | 0.08 | 0.09 | 0.08 | 0.09 | 0.08 | 0.08 | 0.06 |
| **Panel B. Education and employment** | | | | | | | |
| Completed college | 0.62 | 0.62 | 0.62 | 0.65 | 0.60 | 0.67 | 0.58 |
| Numeracy index | 4.09 | 4.00 | 4.12 | 4.16 | 4.11 | 4.15 | 4.03 |
| Works full-time | 0.58 | 0.62 | 0.53 | 0.64 | 0.55 | 0.56 | 0.58 |
| Works part-time | 0.15 | 0.15 | 0.14 | 0.12 | 0.18 | 0.14 | 0.15 |
| Retired | 0.20 | 0.16 | 0.22 | 0.19 | 0.19 | 0.23 | 0.22 |
| Houseowner | 0.72 | 0.68 | 0.77 | 0.70 | 0.69 | 0.74 | 0.77 |
| **Panel C. Household Income** | | | | | | | |
| Income below 40k | 0.24 | 0.25 | 0.22 | 0.19 | 0.23 | 0.24 | 0.27 |
| Income btw. 40k and 75k | 0.28 | 0.30 | 0.27 | 0.28 | 0.28 | 0.25 | 0.26 |
| Income btw. 75k and 100k | 0.16 | 0.17 | 0.14 | 0.16 | 0.17 | 0.15 | 0.17 |
| Income over 100k | 0.32 | 0.25 | 0.36 | 0.38 | 0.31 | 0.36 | 0.28 |
| **Panel D. Region** | | | | | | | |
| Lives in the Midwest | 0.23 | 0.30 | 0.20 | 0.26 | 0.20 | 0.20 | 0.20 |
| Lives in the Northeast | 0.21 | 0.20 | 0.20 | 0.19 | 0.23 | 0.22 | 0.23 |
| Lives in the South | 0.35 | 0.32 | 0.30 | 0.36 | 0.31 | 0.36 | 0.43 |
| Lives in the West | 0.22 | 0.18 | 0.30 | 0.20 | 0.25 | 0.22 | 0.15 |
| Used a mobile device | 0.37 | 0.38 | 0.39 | 0.35 | 0.42 | 0.31 | 0.38 |
| Minutes spent in experiment | 14.72 | 14.67 | 15.06 | 13.62 | 14.83 | 15.37 | 15.09 |
| Passed attention check | 0.89 | 0.86 | 0.85 | 0.92 | 0.93 | 0.91 | 0.88 |
| $N$ (Spectators) | 1,170 | 197 | 194 | 193 | 197 | 196 | 193 |
| $N$ (Redistributive decisions) | 14,040 | 2,364 | 2,328 | 2,316 | 2,364 | 2,352 | 2,316 |

**Notes**: This table shows the demographic composition of our spectator sample, comparing spectators treated with and without information about the likelihood that the best performer is the winner, between lucky outcomes (columns 2 and 5), lucky opportunities (columns 3 and 6), and ex-post lucky opportunities (columns 4 and 7) conditions.

Table 2: Fraction redistributed as a function of the probability that the best performer lost $(1-\pi)$

| | Outcome: Fraction of earnings redistributed | | |
|---|---|---|---|
| | Lucky Outcomes (1) | Lucky Opportunities (2) | Difference (3) |
| **Panel A. Average redistribution** | | | |
| Constant | $0.276^{***}$ | $0.234^{***}$ | $0.042^{**}$ |
| | (0.010) | (0.013) | (0.016) |
| **Panel B. Elasticity of redistribution** | | | |
| Prob. best performer lost $(1-\pi)$ | $0.037^{***}$ | $0.020^{***}$ | $0.017^{**}$ |
| | (0.005) | (0.004) | (0.006) |
| Constant | $0.331^{***}$ | $0.263^{***}$ | $0.068^{***}$ |
| | (0.013) | (0.014) | (0.019) |
| **Panel C. Average redistribution across $\pi$ bins** | | | |
| $1-\pi = 0.00$ | $0.131^{***}$ | $0.179^{***}$ | $-0.048$ |
| | (0.018) | (0.018) | (0.025) |
| $1-\pi \in (0.00, 0.05]$ | $0.175^{***}$ | $0.198^{***}$ | $-0.024$ |
| | (0.018) | (0.016) | (0.024) |
| $1-\pi \in (0.05, 0.10]$ | $0.202^{***}$ | $0.202^{***}$ | $-0.000$ |
| | (0.016) | (0.016) | (0.022) |
| $1-\pi \in (0.10, 0.15]$ | $0.240^{***}$ | $0.226^{***}$ | 0.014 |
| | (0.017) | (0.017) | (0.023) |
| $1-\pi \in (0.15, 0.20]$ | $0.270^{***}$ | $0.208^{***}$ | $0.062^{**}$ |
| | (0.017) | (0.015) | (0.023) |
| $1-\pi \in (0.20, 0.25]$ | $0.316^{***}$ | $0.226^{***}$ | $0.090^{***}$ |
| | (0.016) | (0.015) | (0.022) |
| $1-\pi \in (0.25, 0.30]$ | $0.345^{***}$ | $0.228^{***}$ | $0.117^{***}$ |
| | (0.017) | (0.016) | (0.023) |
| $1-\pi \in (0.30, 0.35]$ | $0.322^{***}$ | $0.240^{***}$ | $0.082^{***}$ |
| | (0.016) | (0.016) | (0.023) |
| $1-\pi \in (0.35, 0.40]$ | $0.315^{***}$ | $0.249^{***}$ | $0.066^{**}$ |
| | (0.016) | (0.016) | (0.023) |
| $1-\pi \in (0.40, 0.45]$ | $0.336^{***}$ | $0.260^{***}$ | $0.076^{**}$ |
| | (0.018) | (0.017) | (0.025) |
| $1-\pi \in (0.45, 0.50)$ | $0.327^{***}$ | $0.291^{***}$ | 0.037 |
| | (0.019) | (0.018) | (0.026) |
| $1-\pi = 0.50$ | $0.336^{***}$ | $0.298^{***}$ | 0.038 |
| | (0.020) | (0.019) | (0.028) |
| $N$ (Redistributive decisions) | 2,364 | 2,328 | 4,692 |

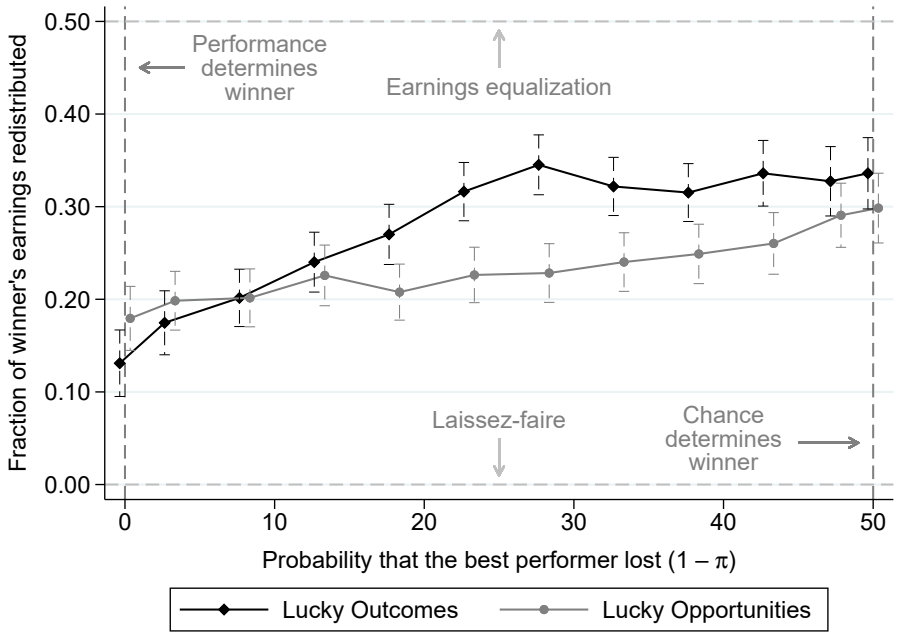
**Notes:** This table shows estimates of redistribution under lucky outcomes (column 1), lucky opportunities (column 2), and the difference (column 3). Panel A shows the mean share of earnings redistributed. Panel B shows the elasticity of the fraction of earnings redistributed with respect to the likelihood that the winning worker performed better than the losing worker. Panel C shows the relationship between redistribution and the likelihood that the winning worker performed better split into 12 bins. Heteroskedasticity-robust standard errors clustered at the spectator level in parentheses. $^{***}$, $^{**}$ and $^{*}$ denote significance at the 0.1%, 1%, and 5% level, respectively.

Table 3: The impact of unequal opportunities on redistribution decisions

| | Outcome: Fraction of earnings redistributed | | |
|---|---|---|---|
| | (1) | (2) | (3) |
| Multiplier difference | 0.041*** | | 0.051*** |
| | (0.007) | | (0.015) |
| Multiplier ratio | | 0.042*** | −0.012 |
| | | (0.008) | (0.015) |
| Constant | 0.202*** | 0.171*** | 0.213*** |
| | (0.014) | (0.017) | (0.019) |
| $N$ (Redistributive decisions) | 2,328 | 2,328 | 2,328 |

**Notes:** This table shows the average redistribution (from the winner's earnings to the loser) as a function of the difference between the multiplier of the winner and the loser ($m_1 - m_2$) and the ratio between the multiplier of the winner and the multiplier of the loser ($m_1/m_2$) in the lucky opportunities environment. Heteroskedasticity-robust standard errors clustered at the spectator level in parentheses. ***, ** and * denote significance at the 0.1%, 1%, and 5% level, respectively.

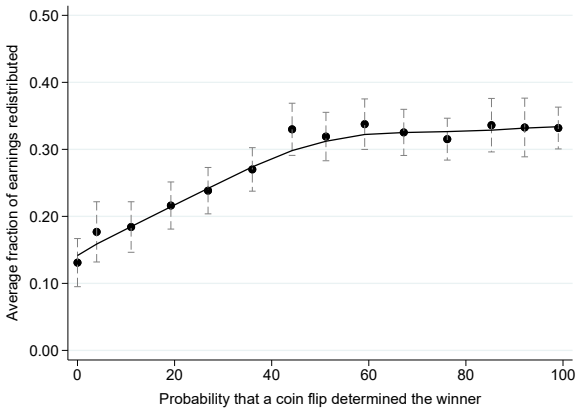
Figure 1: Redistribution and the probability that the best performer lost $(1 - \pi_b)$



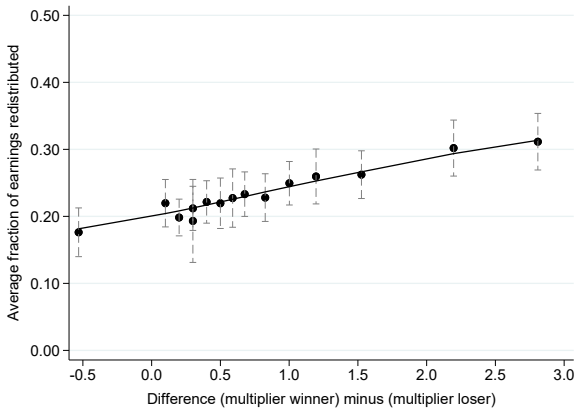

**Notes:** This figure shows the average share of earnings redistributed between workers (from the higher-earning winner to the lower-earning loser) relative to the likelihood that the best performer lost $(1 - \pi)$. The two series represent our main experimental conditions: lucky outcomes (black) and lucky opportunities (grey). The vertical dashed lines denote 95% confidence intervals.

Figure 2: Relationship between the redistribution and the features of the luck environment observed by spectators

Panel A. Average redistribution vs. Coin-flip probability



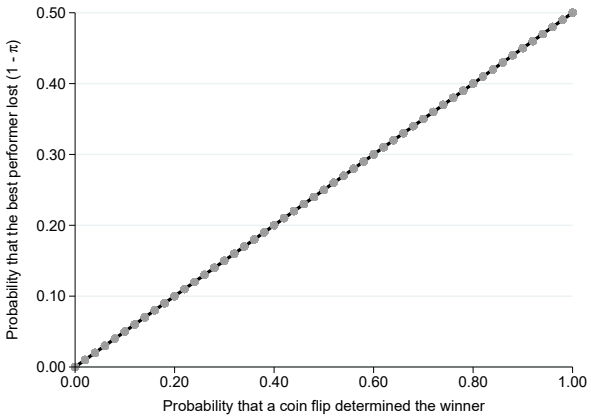

Panel B. Average redistribution vs. multiplier differences



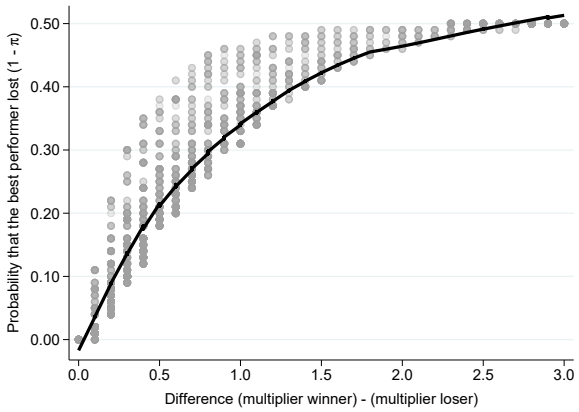

**Notes**: This figure shows the average earnings redistributed from the winner to the loser ($y$-axis) as a function of the probability that a coin flip determined the winner (Panel A) or the difference between the productivity multiplier of the winner and the loser (Panel B), both divided into 15 equal-sized bins. Solid lines plot the predicted values from local polynomial estimates. Dashed vertical lines show 95% confidence intervals,

Figure 3: Relationship between the probability that the best performer lost and features of the luck environment observed by spectators

Panel A. Likelihood that the best performer lost vs. coin-flip probability

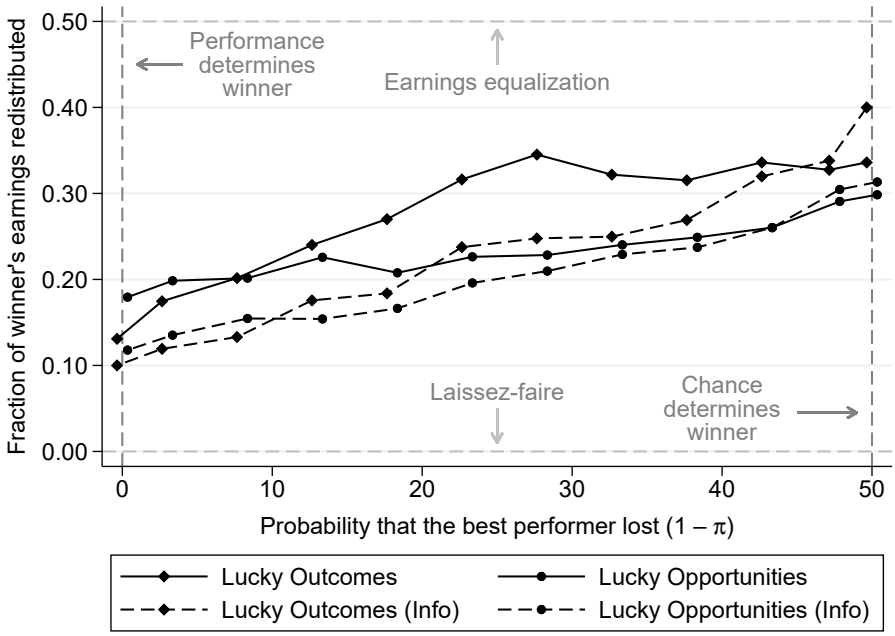



Panel B. Likelihood that the best performer lost vs. multiplier difference



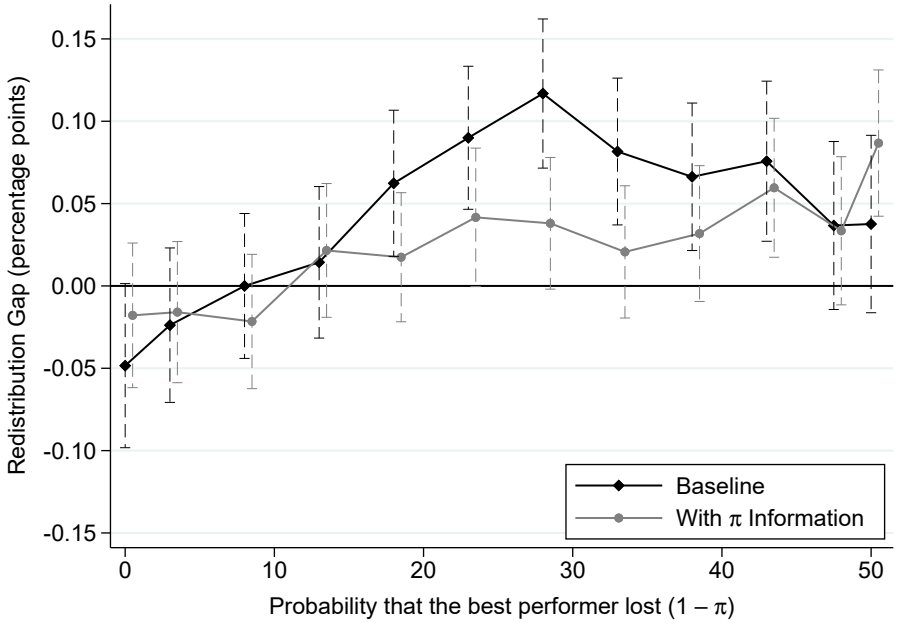

**Notes**: This figure shows the probability that the best performer lost ($y$-axis) as a function of the probability that a coin flip determined the winner (Panel A) or the difference between the productivity multiplier of the winner and the loser (Panel B). The dots denote the empirical frequencies in our experimental data and the intensity of the grey color indicates the number of observations at a given point. The solid lines depict the best fit using either linear specification (panel A) or lowess regression (panel B).

Figure 4: The effect of providing information about $\pi$




**Notes:** This figure shows the average share of earnings redistributed between workers (from the higher-earning winner to the lower-earning loser) as a function of the likelihood that the best performer lost for our two luck environments with and without information about $\pi$. The solid lines depict our baseline lucky outcomes (diamond markers) and lucky opportunities environments (circle markers). The dashed lines depict redistribution for the lucky outcomes (diamond markers) and lucky opportunities (circle markers) environments with information about $\pi$.

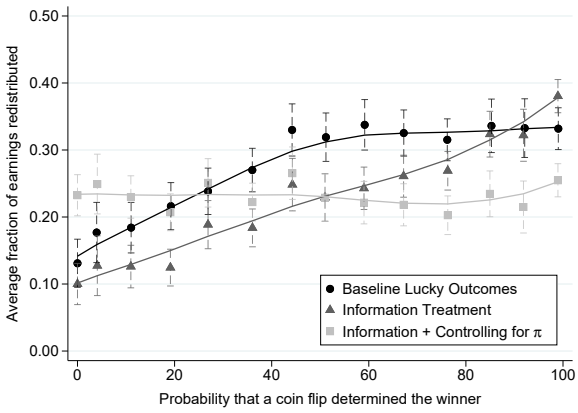
Figure 5: Redistribution gap by $1 - \pi$ in the baseline and information treatments



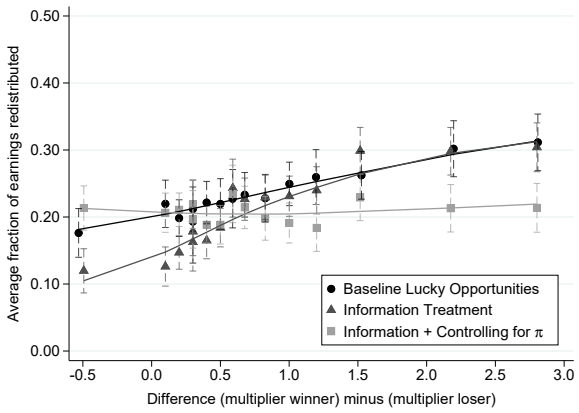

**Notes:** This figure shows the gap in the average share of earnings redistributed between workers (from the higher-earning winner to the lower-earning loser) between the lucky outcomes and lucky opportunities environments as a function of the likelihood that the best performer lost. The solid black line depicts the baseline treatments in which spectators do not observe $\pi$. The solid grey line depicts the information treatment in which spectators have direct access to the exact value of $\pi$ when making their redistribution decisions.

Figure 6: Relationship between the redistribution and observable luck features (information intervention)

Panel A. Redistribution vs. Coin-flip probability

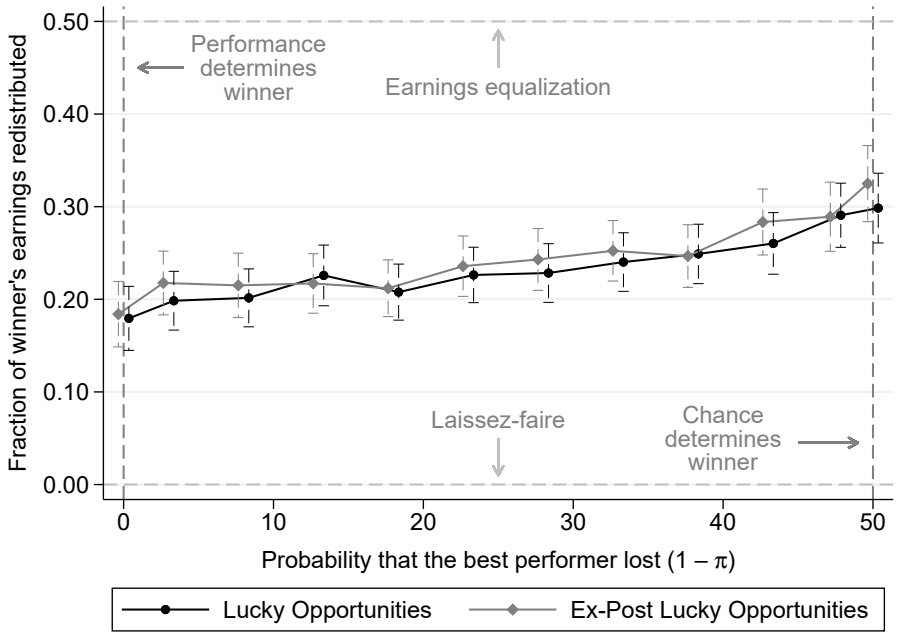



Panel B. Redistribution vs. Multiplier Differences




**Notes**: Panel A shows the average share of winners' earnings redistributed by the probability that a coin flip determined the outcomes. Panel B shows the average share redistributed by the gap in worker multipliers. Information-treated spectators are provided information about $\pi$ in addition to the observable features of the luck environment.

Figure 7: Redistribution by $1 - \pi$ in the baseline and ex-post lucky opportunities conditions




**Notes:** This figure shows the average share of earnings redistributed between workers (from the higher-earning winner to the lower-earning loser) as a function of the likelihood that the best performer lost. It depicts two variations of the lucky opportunities condition: Workers in the baseline condition are aware of their multiplier prior to beginning the encryption task, and workers in the ex-post condition learn their multiplier after completing the encryption task.