

Experimental Design: Main experiment

Participants worked on a real-effort task for three stages (see Figure 1 below); they received piece rate incentives in Stage 1, tournament incentives in Stage 2, and could choose among the two previous incentives in Stage 3. One of the three stages was randomly selected for payment. There were three between-subject treatments, which differed in the nature of the competition in Stage 2 –“meritocratic” or “unfair”– and the type of feedback provided in that stage. We also included a questionnaire eliciting about 50 potential socio-economic determinants of the response to winning or losing. Our pre-analysis plan and power calculation are available from <https://www.socialscienceregistry.org/trials/7651>.

We conducted the experiment on Prolific in Spring 2021. Participants took a median time of 13 minutes and 35 seconds to complete the study. Upon completion, they received a participation fee of £1.50 and a bonus payment based on their decisions and task performance. Feedback on their earnings was provided after the questionnaire. The median payment, including the show-up fee and task earnings, amounted to £2.10.

Meritocratic Treatment.– The real-effort task we used consisted of counting the number of zeros in 8x8 tables consisting of zeros (0) and ones (1). Participants had 90 seconds to solve as many tables as they could, where solving a table meant reporting the correct number of zeros in the respective table. We will refer to the number of solved tables as their “score”. We deviated from the task used in most previous laboratory experiments (addition problems) to reduce the potential for cheating in the online setting, e.g., by using a calculator.

In Stage 1, participants were paid a piece rate of 0.15 pounds per table they solved correctly. In Stage 2, participants were instead remunerated according to a two-person winner-takes-all tournament. Participants were told that they would receive 0.30 pounds per table if their score exceeded the score of a random opponent who had already completed the task, and zero otherwise. Ties were broken randomly. The opponent’s score was randomly selected from the score distribution of participants in Apicella et al. 2017. In Stage 3, participants made a choice whether to apply piece rate or tournament pay to their performance. If they chose the tournament, their opponent’s score was once again randomly selected from the same score distribution of prior participants.

Apart from the task, our design differed from Niederle & Vesterlund 2007 in

Stage 1: Piece Rate

- Solve tables for 90 seconds with piece rate incentives (£0.15 per table)
- Feedback: “You scored ... correct answers.”

Stage 2: Tournament

- Solve tables for 90 seconds with tournament incentives (£0.30 per table for winners, 0 for losers)
 - ★ *Meritocratic Treatment.*– Best performer wins, ties broken randomly
 - ★ *Unfair/Feedback Treatment.*– Best performer wins 75% of the time, worst performer wins 25% of the time, ties broken randomly
- Prior belief elicitation: tournament rank in Stage 2 (£0.50-£0.02*guessing error)
- Feedback: “You scored ... correct answers. You won/lost in the tournament in Stage 2.”
 - ★ *Feedback Treatment:* Additional feedback consisting of performance rank relative to the comparison sample + information on whether their score exceeded the score of the opponent and they (un)deservedly won/lost the tournament.
- *Meritocratic/Unfair Treatment.*– Posterior belief elicitation: tournament rank in Stage 2 (£0.50-£0.02*guessing error)

Stage 3: Choice

- Choose between piece rate and tournament incentives
 - ★ *All Treatments.*– Best performer wins, ties broken randomly
- Solve tables for 90 seconds under chosen incentive
- Feedback: “You scored ... correct answers.” (if piece rate was chosen)
- Feedback: “You scored ... correct answers. You won/lost in the tournament in Stage 3.” (if tournament was chosen)

Questionnaire

- Demographics + preferences for risk and competition

Payment Screen

- One Stage randomly selected for payment

Figure 1: Overview of the Experiment

two main ways. First, we informed participants after Stage 2 whether they won or lost the tournament. This difference is crucial to study the effect of failure (losing a competition) on tournament entry in Stage 3. Second, we included two incentivised belief elicitation tasks in Stage 2, one before and one after receiving information on whether they won or lost the tournament. In both elicitation tasks, we asked participants to estimate their rank compared to 100 prior participants in Stage 2. To keep things simple, participants were paid according to a linear scoring rule that awarded them a base payment of 0.50 pounds minus 0.02 pounds times the absolute difference between the true rank and the stated (guessed) rank, with a minimum of zero. In case of ties, the true rank we assigned to participants was equal to the average rank of all prior participants with the same score.

Unfair Treatment.— Our goal in designing our second treatment was to introduce a source of meritocratic unfairness. We use an approach in which distortions away from meritocracy are generated by a random process not linked to gender or other specific demographic characteristics. We refer to this treatment as the “Unfair” Treatment.

In practice, the only difference between this treatment and the Meritocratic Treatment lies in the way the winner was determined in Stage 2. In contrast to the Meritocratic Treatment, where the best performer would always win the tournament, participants in the Unfair Treatment were informed that there was a 75% chance that the best performer would win the tournament. In the remaining 25% of cases, the inferior performer would win instead. Participants were told who won the tournament, but not whether the winner also had the best performance. It is important to note that this change only applied to Stage 2. In Stage 3, the tournament (if chosen) was still a fully meritocratic tournament, where the best performer always won. Prior to Stage 3, participants in all treatments were explicitly informed that the best performer would always win in the tournament (if chosen). This design feature allows us to attribute differences in tournament entry in Stage 3 to the experience of winning or losing a prior competition, as opposed to the willingness to enter fair versus unfair competitions.

Feedback Treatment.— Our final treatment is identical to the Unfair Treatment, except that participants received two additional types of feedback at the end of Stage 2. First, we informed participants about their true rank in Stage 2 relative to the comparison sample of 100 participants. Second, we told participants whether their score was higher than their opponent’s and whether they deservedly (75%

of the time) or undeservedly (25% of the time) won or lost the tournament. This treatment therefore gives participants perfect information about both the nature of the competition and their ability rank. The only other change was that we removed the belief elicitation task at the end of Stage 2, because participants knew their exact rank.

Experimental Design: Survey 1

We presented 670 novel participants with a summary of Stage 2 of the Meritocratic Treatment from our main experiment. We then asked participants whether the process determining the winner was (a) fair, (b) merit-based, and (c) random, each on a scale from 0-100. Participants could rate the treatments using a slider. The three questions were presented in a random order. We then presented participants with a summary of the Unfair Treatment, and asked them the same three questions in the same way. We always present the two treatments in the same order to facilitate comprehension. We closed with an open question that asked participants to explain their reasoning. We did not ask any demographics questions, but have access to the demographic data collected by Prolific itself, including gender, age and country of residence.

Experimental Design: Survey 2

The second survey had a similar structure to the first. In particular, participants were introduced to the Meritocratic Treatment and asked to rate it in the same way as in the first survey. The difference is that following the Meritocratic Treatment they were introduced to a “risk-only” task rather than the Unfair Treatment. In particular, we explained that in the risk-only task, the winner would be determined by a coin-flip rather than through a tournament. We compared this task to the Meritocratic Treatment rather than the Unfair Treatment to facilitate participant comprehension. This survey allows us to compare fairness perceptions between our main experiment and the important benchmark of a random lottery. In particular, it allows us to see whether the Unfair Treatment is seen as similar to a lottery, or differs from a lottery in important ways.