

An AI system trained to find an equitable policy for distributing public funds in an online game

July 5 2022, by Bob Yirka

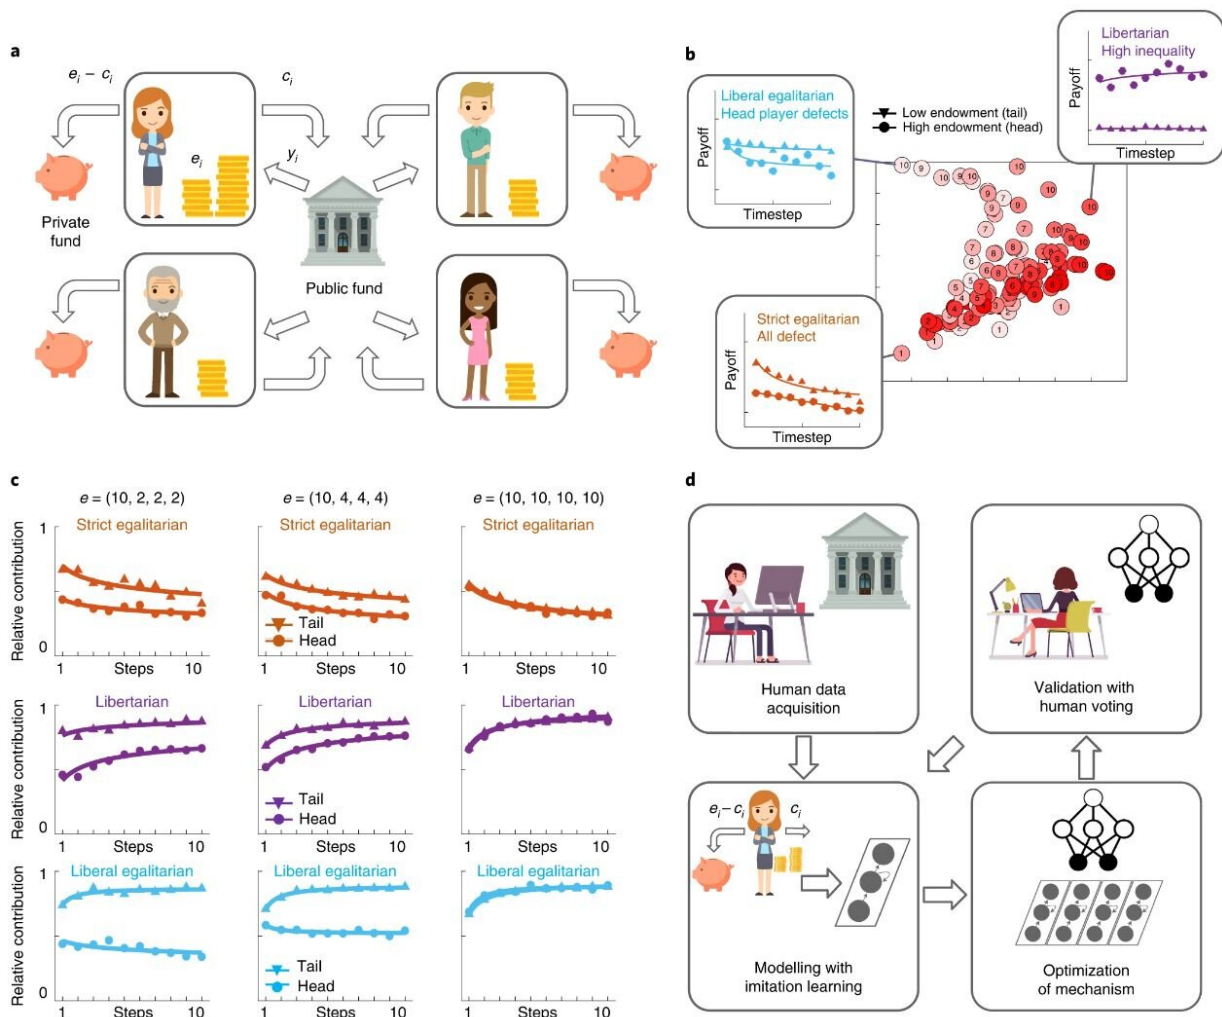


Illustration of the game and Experiment 1. a, Illustration of the setup of the investment game. b, The ideological manifold for endowment distribution (10, 2,

2, 2). The plot shows a visualization of a space of redistribution mechanisms defined by parameters w and v in two dimensions. Each red dot is a mechanism, and distances between dots conserve dissimilarities in the (average) relative payout to virtual players (both head and tail). Dot numbers denote bins of mechanism parameter w (1, lowest; 10, highest) and shading denotes bins of v (light, more relative; dark, more absolute). Inset, example payouts to head (circles) and tail (triangles) players under the canonical mechanisms used as baselines against which to test the AI. Under strict egalitarian, payouts decline to head and tail players. Under libertarian, there is great inequality between head and tail players. Under liberal egalitarian, the head player stops contributing, so payouts decline for both head and tail players. c, Average relative contributions (as a fraction of endowment) over 10 rounds (x axis) in Exp. 1 for three different initial endowment conditions. Under strict egalitarian redistribution, tail player (triangles) contributions are higher when initial endowments are lower, but head player (circles) contributions do not differ. Under libertarian, head player contributions increase with equality, but tail player contributions remain constant. Head player contributions increase strongly with endowment under liberal egalitarian. d, Illustration of our agent design pipeline. Credit: *Nature Human Behaviour* (2022). DOI: 10.1038/s41562-022-01383-x

A team of researchers at DeepMind, London, working with colleagues from the University of Exeter, University College London and the University of Oxford, has trained an AI system to find a policy for equitably distributing public funds in an online game. In their paper published in the journal *Nature Human Behavior*, the group describes the approach they took to training their system and discuss issues that were raised in their endeavor.

How a society distributes wealth is an issue that humans have had to face for thousands of years. Nonetheless, most economists would agree that no system has yet been established in which all of its members are happy with the status quo. There have always been inequitable levels of income, with those on top the most satisfied and those on the bottom the

least satisfied. In this latest effort, the researchers in England took a new approach to solving the problem—asking a computer to take a more logical approach.

The researchers began with the assumption that [democratic societies](#), despite their flaws, are thus far the most agreeable of those tried. They then enlisted the assistance of volunteers to play a simple resource allocation [game](#)—the players of the game decided together the best ways to share their mutual resources. To make it more realistic, the players received different amounts of resources at the outset and there were different distribution schemes to choose from. The researchers ran the game multiple times with different groups of volunteers. They then used the data from all of the games played to train several AI systems on the ways that humans work together to find a solution to such a problem. Next, they had the AI systems play a similar game against one another, allowing for tweaking and learning over multiple iterations.

The researchers found that the AI systems had settled on a form of liberal egalitarianism in which players received few resources unless they contributed proportionally heavily to the community pool. The researchers then finished their research by asking a group of human volunteers to play the same game as before, only this time, they were given a choice between using one of several conventional sharing approaches or the one developed by the AI system—the one devised by the AI system was the consistent choice among the human players.

More information: Raphael Koster et al, Human-centred mechanism design with Democratic AI, *Nature Human Behaviour* (2022). [DOI: 10.1038/s41562-022-01383-x](https://doi.org/10.1038/s41562-022-01383-x)

Citation: An AI system trained to find an equitable policy for distributing public funds in an online game (2022, July 5) retrieved 17 April 2024 from <https://techxplore.com/news/2022-07-ai-equitable-policy-funds-online.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.