LETTER

Social dilemmas among unequals

Oliver P. Hauser[1,6]*, Christian Hilbe[2,6]*, Krishnendu Chatterjee[2] & Martin A. Nowak[3,4,5]*

Direct reciprocity is a powerful mechanism for the evolution of cooperation on the basis of repeated interactions[1–4]. It requires that interacting individuals are sufficiently equal, such that everyone faces similar consequences when they cooperate or defect. Yet inequality is ubiquitous among humans[5,6] and is generally considered to undermine cooperation and welfare[7–10]. Most previous models of reciprocity do not include inequality[11–15]. These models assume that individuals are the same in all relevant aspects. Here we introduce a general framework to study direct reciprocity among unequal individuals. Our model allows for multiple sources of inequality. Subjects can differ in their endowments, their productivities and in how much they benefit from public goods. We find that extreme inequality prevents cooperation. But if subjects differ in productivity, some endowment inequality can be necessary for cooperation to prevail. Our mathematical predictions are supported by a behavioural experiment in which we vary the endowments and productivities of the subjects. We observe that overall welfare is maximized when the two sources of heterogeneity are aligned, such that more productive individuals receive higher endowments. By contrast, when endowments and productivities are misaligned, cooperation quickly breaks down. Our findings have implications for policy-makers concerned with equity, efficiency and the provisioning of public goods.

In social dilemmas, overall welfare is maximized if all individuals cooperate yet each individual prefers to defect[16]. Such dilemmas occur at all levels of human society. They affect families, companies and nations[17,18]. An extensive body of research has shown that cooperation is more likely when groups are stable and subjects interact repeatedly[11–15]. However, this mechanism of direct reciprocity assumes that group members have sufficient leverage to influence one another. Subjects need to be able to give appropriate responses. Tit-for-tat can only be effective if it incentivizes others to cooperate. Most previous models of reciprocity assume perfect symmetry between individuals[11–15]. Real groups often exhibit substantial heterogeneity, which is derived from multiple sources[5,6]. Experimental studies have shown that inequality in the endowments of the players reduces cooperation[7,8] and undermines the social structure of a population[9]. Even if subjects start out equally, game dynamics can introduce inequality over time, disfavouring individuals who are more cooperative[19] (Supplementary Information). So far, it has been difficult to predict the effect of heterogeneity on cooperation, especially if subjects vary along multiple dimensions. Here we propose a general framework to explore how different kinds of heterogeneities interact and affect cooperation.

We consider public goods games with n players. In each round, player i receives a fixed endowment e_i, which can be interpreted as a regular income. After receiving their endowments, players independently decide which fraction x_i of their endowment to contribute to the public good. The payoff u_i of player i for that round depends on the distribution of endowments, $e_1, \ldots, e_n$, and on the relative contributions of the players, $x_1, \ldots, x_n$. It is typically assumed that all players have the same endowment and that contributions to the public goods are multiplied by a common productivity factor, r (Fig. 1a). Here we instead consider

interactions in which players have different endowments, different productivities or for which payoffs are nonlinear (Fig. 1b–d).

As a specific example, we consider public goods games in which the contributions of each player are multiplied by an individual factor r_i and equally shared among all participants,

$$u_i = \frac{1}{n} \sum_{j=1}^{n} r_j e_j x_j + (1 - x_i) e_i \quad (1)$$

The first term represents the payoff that is derived from the public good and the second term represents the remaining endowment of the player. We interpret the factors r_j as the productivity of each player and assume $1 < r_j < n$ for all j. Thus, the game is a social dilemma in which individuals have an incentive to free-ride[16]. Although we focus on the example given by equation (1) throughout most of the main text, our findings generalize to arbitrary public goods games that satisfy four natural requirements (Methods).

If the public goods game is played once, defection is the only equilibrium. But for repeated interactions, cooperation can prevail if players adopt conditional strategies, such as tit-for-tat[2] or win-stay lose-shift (WSLS)[11], or multiplayer variants of these strategies[20]. We assume after each round that there is another round with probability δ.

To explore the effects of different kinds of heterogeneity, we first characterize when it is that cooperation can be maintained. For a
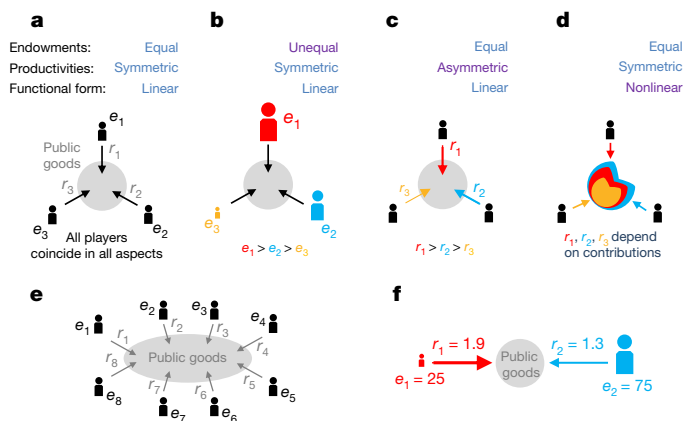

Fig. 1 | Public goods games among unequals. We consider social dilemmas in which participants decide how much of their endowment e_i to contribute to the public goods. The contributions of each player are multiplied by r_i and then divided among all players. The players have equal endowments if $e_1 = e_2 = e_3$. The game is symmetric if players are indistinguishable except for their endowments and contributions. Here the game is symmetric if $r_1 = r_2 = r_3$. The game is linear if the payoffs depend linearly on the endowments and contributions of the players. Here the game is linear if the factors r_i are constant. **a**, Most previous studies assume that players have the same endowment, the game is symmetric and payoffs are linear. **b–d**, Instead, we allow players to have different endowments, different productivities and nonlinear payoffs. **e**, We derive general results for n-player games. **f**, As a special case, we study pairwise interactions.

[1]Department of Economics, University of Exeter Business School, Exeter, UK. [2]IST Austria, Klosterneuburg, Austria. [3]Program for Evolutionary Dynamics, Harvard University, Cambridge, MA, USA. [4]Department of Organismic and Evolutionary Biology, Harvard University, Cambridge, MA, USA. [5]Department of Mathematics, Harvard University, Cambridge, MA, USA. [6]These authors contributed equally: Oliver P. Hauser, Christian Hilbe. *e-mail: o.hauser@exeter.ac.uk; christian.hilbe@ist.ac.at; martin_nowak@harvard.edu
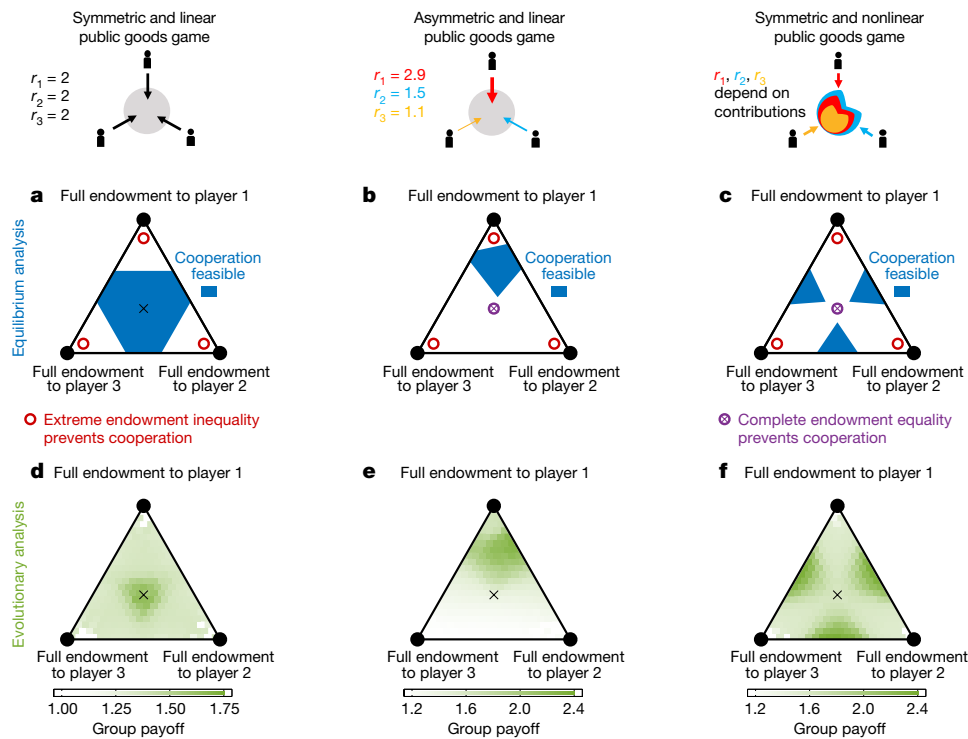
Fig. 2 | Feasibility and evolvability of cooperation in public goods games among unequals. a–f, We consider groups of three players who interact in three different public goods games. In each case, we investigated when equal endowments help to maintain cooperation (**a–c**) or favour its evolution (**d–f**). The triangles represent the possible ways to distribute the initial endowment among the players. Corners correspond to distributions for which one player receives all of the endowment. Edges correspond to distributions for which one player receives no endowment. The centre of the triangle marks equal endowments. 'Group payoff' corresponds to the total payoff across all group members, averaged over 10^6 time steps of an evolutionary simulation. We find that extreme inequality is always detrimental to cooperation. However, when the game is asymmetric or nonlinear, slightly unequal endowments may be necessary for cooperation to be feasible (**b, c**) and for cooperation to evolve (**e, f**).

given public goods game and a given endowment distribution, we say that full cooperation is feasible if there is a subgame perfect equilibrium in which all players always contribute their entire endowment. In such an equilibrium, players have no incentive to deviate after any history of previous play[21]. In the Supplementary Information, we prove that cooperation is feasible if and only if the strategy Grim is an equilibrium. Grim cooperates unless another player has defected in a previous round[3]. From the equilibrium condition for Grim, it follows that cooperation is feasible for the public good game given by equation (1) if and only if for all players i with $e_i > 0$ the following condition holds:

$$\frac{\delta}{n} \sum_{j \neq i} r_j e_j \geq \left(1 - \frac{r_i}{n}\right) e_i \qquad (2)$$

The expected benefit from the future cooperation of others must exceed the incentive to defect in the present round. For cooperation to be feasible, future losses must outweigh present gains.

On the basis of this general characterization of when cooperation is feasible, we derive a number of results. First, cooperation is never feasible if there is too much inequality, such that most of the endowment is in the hands of one player (Supplementary Information). For linear and symmetric games (Fig. 1), we show that if cooperation is feasible at all, it is feasible for equal endowments (Fig. 2a). However, if the game is asymmetric (Fig. 2b) or nonlinear (Fig. 2c), full cooperation may only be feasible when players have unequal endowments. In such a case it can even be optimal to give some players no initial endowment at all.

To gain intuition, consider a case in which players differ in productivities, $r_1 > \ldots > r_n$. We find a twofold advantage of giving higher endowments to more productive players. First, there is a stability advantage: an unequal distribution of endowments makes it easier for full cooperation to be an equilibrium. To understand this, assume instead that

players receive equal endowments. Then inequality (2) suggests that cooperation is feasible if

$$\delta \geq \frac{n - r_n}{R - r_n} \qquad (3)$$

in which $R = r_1 + \ldots + r_n$ is the sum of all productivities. For equal endowments, player n with the lowest productivity faces the largest temptation to defect, because this player has the highest marginal cost $1 - r_n/n$ of contributing. This temptation can be counterbalanced by allocating a smaller endowment to player n who then has less to gain from withholding, whereas the others have more leverage to retaliate in future rounds. Both effects enhance the stability of cooperation. Second, there is an efficiency advantage of unequal endowments. Because contributions of more productive players are multiplied by a higher factor, social welfare is maximized when the most productive player obtains the largest share of the initial endowment—subject to the constraint that full cooperation is feasible.

If the game involves only two players, we can compute which endowment distribution is the most conducive to cooperation. An endowment distribution is maximally cooperative if it requires the lowest continuation probability δ for cooperation to be feasible. Using inequality (2), we show in the Supplementary Information that endowments need to be distributed as

$$\frac{e_1}{e_2} = \sqrt{\frac{r_2(2 - r_2)}{r_1(2 - r_1)}} \qquad (4)$$

An equal distribution, $e_1 = e_2$, is maximally cooperative only if players have the same productivities. Otherwise, the more productive player should have a larger share of the endowment.

After exploring under which conditions cooperation is feasible, we study when it is that cooperation can evolve[3]. To make an evolutionary
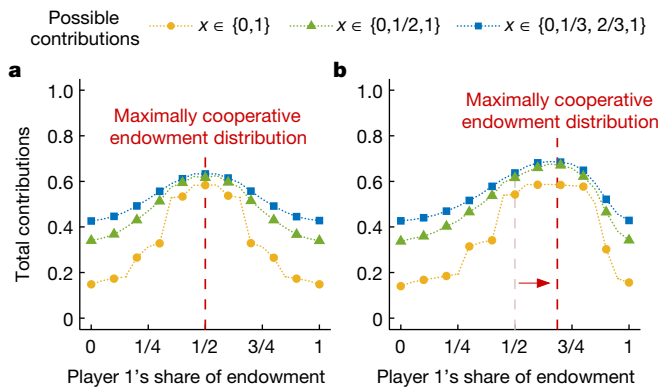
Fig. 3 | When players differ in their productivities, equal endowments do not maximize contributions. We consider public goods games between two players. **a, b,** Players either coincide in their productivities, $r_1 = r_2 = 1.6$ (**a**) or player 1 is more productive, $r_1 = 1.9$ and $r_2 = 1.3$ (**b**). In each case, we vary player 1's share of the initial endowment. We perform evolutionary simulations (dotted lines) for three scenarios, depending on which fraction of the endowment the players can contribute: $\{0, 1\}$, $\{0, 1/2, 1\}$ or $\{0, 1/3, 2/3, 1\}$. For equal productivities, cooperation is most likely to emerge when both players receive the same endowment. By contrast, when players differ in productivity, the maximum total payoff is achieved when player 1 obtains a larger fraction of the endowment. The position of the maximum is well-approximated by the maximally cooperative endowment distribution given by equation (4).

approach computationally tractable, we first consider players who respond only to the outcome of the last round. Moreover, we assume that players choose only from a finite set of possible contributions. For example, they may either contribute their full endowment or nothing at all. In that case, we refer to the two possible actions as cooperation and defection, respectively. With some small probability, ε, players commit errors such that a player who intends to cooperate defects by mistake (and vice versa). Players adopt new strategies over time by comparing their payoff to the payoff they would obtain by using a random alternative strategy. The better the payoff of the alternative strategy, the more likely players are to switch. We iterate this process for many steps and record the average cooperation rates over time (Methods).

Our numerical findings parallel the previous equilibrium results. Cooperation cannot evolve if one of the players receives almost all of the endowment. Moreover, for linear and symmetric games, individuals are most likely to cooperate if everyone receives the same endowment (Fig. 2d). However, if some players are more productive than others (Fig. 2e) or if the game is nonlinear (Fig. 2f), unequal endowments yield more cooperation and higher payoffs. In all cases, we observe that the strategy Grim is less relevant, because it cannot sustain cooperation in the presence of noise[3]. Instead, cooperation evolves if the strategy WSLS[11] is an equilibrium (Extended Data Figs. 1–5). WSLS contributes the full endowment in the first round, or if all players made the same relative contribution in the previous round. Otherwise WSLS contributes nothing[11,20].

In the simulations, a group of defectors is most likely to be invaded by strategies such as tit-for-tat. These conditional cooperators in turn quickly adopt WSLS, which is more robust with respect to errors. However, because of stochasticity, any strategy is replaced eventually, even if it is an equilibrium (Supplementary Information). Further simulations show that analogous results hold when players choose between more than two discrete contributions each round (Fig. 3) or when strategies are represented by finite-state automata[15] (Extended Data Fig. 6, see Supplementary Information for details).

To explore the applicability of these theoretical results, we designed an online behavioural experiment based on the two-player game of Fig. 3. Participants are either equally productive or not and have the same endowment or not. We consider five treatments: full equality, endowment inequality, productivity inequality, aligned inequality and
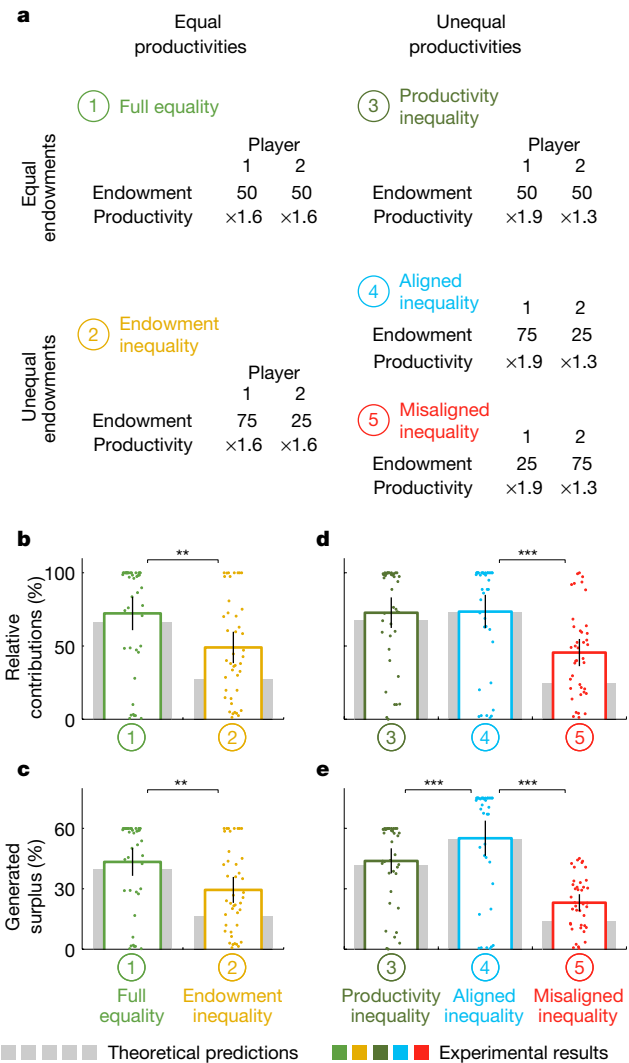

Fig. 4 | Exploring the effects of multidimensional inequality with a behavioural experiment. a, On the basis of the two-player game shown in Fig. 3, we conduct an experiment with varying endowments and productivities. There are five conditions: (1) full equality, (2) endowment inequality, (3) productivity inequality, (4) aligned inequality (the more productive player has higher endowment), and (5) misaligned inequality (the more productive player has lower endowment). **b–e,** For each treatment, we compare the theoretical predictions from evolutionary simulations (grey bars) with the respective average values of the experiment (coloured bars). We show the relative contributions of each player (top) and the generated surplus (by how much the total payoffs of the players exceed their initial endowments; bottom). Aligned inequality yields high cooperation rates and higher payoffs than other treatments. Coloured dots represent individual groups of players; the number of observations (groups) for each treatment was 42, 42, 40, 39, 40 for treatments 1–5, respectively. Error bars represent 95% confidence intervals. We analysed pairwise differences between treatments using two-tailed Mann–Whitney U-tests. $**P < 0.01$; $***P < 0.001$. See Methods for details.

misaligned inequality (Fig. 4a). In the last two treatments, individuals differ in both dimensions. Either the more productive player (aligned) or the less productive player (misaligned) receives the larger endowment. Previous experiments suggest that—in isolation—heterogeneous endowments reduce cooperation[7,8], whereas heterogeneous productivities have a negligible effect[22]. Here we study the interaction between the two heterogeneities in repeated games, for which previous research did not find any significant effects[23].

On the basis of our evolutionary analysis, we expect aligned inequality to increase and misaligned inequality to reduce welfare compared

to the case of productivity inequality alone (Fig. 3). The experiment confirms these predictions (Fig. 4b–e). Aligned inequality results in substantially higher contributions than misaligned inequality and generates the highest surplus across all treatments. Under aligned inequality, most high-endowment players match the relative contribution of the low-endowment players. That is, if the low-endowment player gives their full endowment, then so does the high-endowment player, even if their absolute contributions are three times higher. By contrast, contributions under misaligned inequality do not follow a clear norm; often, the high-endowment player only matches the absolute contribution of the other player (Extended Data Figs. 7–10 and Supplementary Information).

Here we introduce a general framework to study direct reciprocity among unequals. Our three complementary approaches—equilibrium calculations, evolutionary simulations and a behavioural experiment—suggest an unexpected benefit of inequality. We show that equal endowments can be detrimental to social welfare if subjects differ along multiple other dimensions, such as productivity or benefits from public goods. In those cases, some inequality can increase both the stability of cooperation and the efficiency of contributions.

Despite these potential benefits, inequality comes with caveats. First, maximizing cooperation requires a delicate balance between the different dimensions of heterogeneity. Finding the right amount of inequality can prove difficult when the players' personal characteristics, such as their productivities, are known only imperfectly. The problem is aggravated by our finding that an excess of inequality is always detrimental.

Second, endowment inequality could interfere with institutional solutions to cooperation. For example, when cooperation is maintained through sanctions, heterogeneous groups may disagree on which norm to enforce[24]. Additional problems arise when sanctioning institutions can be corrupted[25,26], especially when better-endowed individuals can 'play the system'.

Finally, reducing inequality is often considered an important policy objective in itself. Humans dislike inequality[27] and are sometimes willing to sacrifice their own wealth to guarantee more-egalitarian outcomes[28,29]. In addition, inequality often renders successful coordination in social dilemmas difficult, as different actors may disagree on which cooperative equilibrium is fair[30]. However, here we show that inequality does not need to render cooperation impossible. When individuals are naturally heterogeneous, moderate inequality can be necessary for cooperation to prevail.

Online content

Any methods, additional references, Nature Research reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at https://doi.org/10.1038/s41586-019-1488-5.

1. Trivers, R. L. The evolution of reciprocal altruism. *Q. Rev. Biol.* **46**, 35–57 (1971).
2. Axelrod, R. *The Evolution of Cooperation* (Basic Books, 1984).
3. Sigmund, K. *The Calculus of Selfishness* (Princeton Univ. Press, 2010).
4. Nowak, M. A. Five rules for the evolution of cooperation. *Science* **314**, 1560–1563 (2006).
5. Piketty, T. & Saez, E. Inequality in the long run. *Science* **344**, 838–843 (2014).
6. Scheffer, M., van Bavel, B., van de Leemput, I. A. & van Nes, E. H. Inequality in nature and society. *Proc. Natl Acad. Sci. USA* **114**, 13154–13157 (2017).
7. Cherry, T. L., Kroll, S. & Shogren, J. F. The impact of endowment heterogeneity and origin on public good contributions: evidence from the lab. *J. Econ. Behav. Organ.* **57**, 357–365 (2005).
8. Hargreaves Heap, S. P., Ramalingam, A. & Stoddard, B. Endowment inequality in public good games: a re-examination. *Econ. Lett.* **146**, 4–7 (2016).
9. Nishi, A., Shirado, H., Rand, D. G. & Christakis, N. A. Inequality and visibility of wealth in experimental social networks. *Nature* **526**, 426–429 (2015).
10. Hauser, O. P., Kraft-Todd, G. T., Rand, D. G., Nowak, M. A. & Norton, M. I. Invisible inequality leads to punishing the poor and rewarding the rich. *Behav. Public Policy* https://doi.org/10.1017/bpp.2019.4 (2019).
11. Nowak, M. & Sigmund, K. A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner's Dilemma game. *Nature* **364**, 56–58 (1993).
12. Szabó, G., Antal, T., Szabó, P. & Droz, M. Spatial evolutionary prisoner's dilemma game with three strategies and external constraints. *Phys. Rev. E* **62**, 1095–1103 (2000).
13. Doebeli, M. & Hauert, C. Models of cooperation based on the prisoner's dilemma and the snowdrift game. *Ecol. Lett.* **8**, 748–766 (2005).
14. Stewart, A. J. & Plotkin, J. B. From extortion to generosity, evolution in the Iterated Prisoner's Dilemma. *Proc. Natl Acad. Sci. USA* **110**, 15348–15353 (2013).
15. van Veelen, M., García, J., Rand, D. G. & Nowak, M. A. Direct reciprocity in structured populations. *Proc. Natl Acad. Sci. USA* **109**, 9929–9934 (2012).
16. Kerr, B., Godfrey-Smith, P. & Feldman, M. W. What is altruism? *Trends Ecol. Evol.* **19**, 135–140 (2004).
17. Rand, D. G. & Nowak, M. A. Human cooperation. *Trends Cogn. Sci.* **17**, 413–425 (2013).
18. Frank, M. R. et al. Detecting reciprocity at a global scale. *Sci. Adv.* **4**, eaao5348 (2018).
19. Gächter, S., Mengel, F., Tsakas, E. & Vostroknutov, A. Growth and inequality in public good provision. *J. Public Econ.* **150**, 1–13 (2017).
20. Pinheiro, F. L., Vasconcelos, V. V., Santos, F. C. & Pacheco, J. M. Evolution of all-or-none strategies in repeated public goods dilemmas. *PLOS Comput. Biol.* **10**, e1003945 (2014).
21. Fudenberg, D. & Tirole, J. *Game Theory* 6th edn (MIT Press, 1998).
22. Fisher, J., Isaac, R. M., Schatzberg, J. W. & Walker, J. M. Heterogenous demand for public goods: behavior in the voluntary contributions mechanism. *Public Choice* **85**, 249–266 (1995).
23. van Gerwen, N., Buskens, V. & van der Lippe, T. Individual training and employees' cooperative behavior: evidence from a contextualized laboratory experiment. *Rationality Soc.* **30**, 432–462 (2018).
24. Reuben, E. & Riedl, A. Enforcement of contribution norms in public good games with heterogeneous populations. *Games Econ. Behav.* **77**, 122–137 (2013).
25. Abdallah, S. et al. Corruption drives the emergence of civil society. *J. R. Soc. Interface* **11**, 20131044 (2014).
26. Muthukrishna, M., Francois, P., Pourahmadi, S. & Henrich, J. Corrupting cooperation and how anti-corruption strategies may backfire. *Nat. Hum. Behav.* **1**, 0138 (2017).
27. Tricomi, E., Rangel, A., Camerer, C. F. & O'Doherty, J. P. Neural evidence for inequality-averse social preferences. *Nature* **463**, 1089–1091 (2010).
28. Dawes, C. T., Fowler, J. H., Johnson, T., McElreath, R. & Smirnov, O. Egalitarian motives in humans. *Nature* **446**, 794–796 (2007).
29. Durante, R., Putterman, L. & Van der Weele, J. Preferences for redistribution and perception of fairness: an experimental study. *J. Eur. Econ. Assoc.* **12**, 1059–1086 (2014).
30. Tavoni, A., Dannenberg, A., Kallis, G. & Löschel, A. Inequality, communication, and the avoidance of disastrous climate change in a public goods game. *Proc. Natl Acad. Sci. USA* **108**, 11825–11829 (2011).

METHODS

General modelling framework. In the main text, we used the public goods game in equation (1) to illustrate our main findings. However, the framework that we use to study reciprocity in asymmetric social dilemmas is general and can encompass many other examples (as indicated in Fig. 2c, f). Here we introduce, in brief, our general framework. A full account is provided in the Supplementary Information.

We consider games with n players. The endowments of the players in each round are given by the endowment vector $\boldsymbol{e} = (e_1, \ldots, e_n)$. Endowments are non-negative, $e_i \geq 0$ for all players i, and normalized, $e_1 + \ldots + e_n = 1$. Given their endowments, players decide which fraction of their endowment they contribute, as summarized by the contribution vector $\boldsymbol{x} = (x_1, \ldots, x_n)$ for which $x_i \in [0, 1]$ for all players i. We refer to x_i as the relative contribution of the player and to $e_i x_i$ as the absolute contribution of the player. We use the shorthand notation $\boldsymbol{x} = \boldsymbol{0}$ if no player contributes to the public goods and $\boldsymbol{x} = \boldsymbol{1}$ if all players contribute their full endowment. Given the endowments $\boldsymbol{e}$ and the relative contributions $\boldsymbol{x}$ in a given round, the payoff for player i in that round is $u_i(\boldsymbol{e}, \boldsymbol{x})$. If there are no contributions, players receive their initial endowment $u_i(\boldsymbol{e}, \boldsymbol{0}) = e_i$.

We consider public goods games that satisfy the following four conditions: continuity (C), the payoff functions $u_i(\boldsymbol{e}, \boldsymbol{x})$ are continuous in both arguments; positive externalities (PE), as a player with a positive endowment increases their contribution, the payoffs of all other players increase; incentive to free-ride (IF), as a player with positive endowment increases their contribution, their own payoff decreases; and optimality of cooperation (OC), as a player with positive endowment increases their contribution, the overall payoff over all members of the group increases.

The first condition of continuity is merely a technical assumption that is useful for some of the analytical results. The other three conditions generalize previous notions of social dilemmas in one-shot games in which players can either cooperate or defect[16].

We note that the above conditions rule out certain threshold public goods games, in which payoffs increase discontinuously once total contributions exceed a certain threshold[31–34]. In such threshold public goods games, cooperation can often emerge even if the game is played only once, because players have an incentive not to fall below the threshold[35]. By considering public goods games that satisfy (C), (PE), (IF) and (OC), we consider the most stringent case of a social dilemma, in which repeated interactions are key to sustain positive contributions. Asymmetric threshold public goods games have been studied in a one-shot or finite-horizon setting have previously been studied[36–39]. Previous work has also explored the consequences of heterogeneities in the background fitness of the players[40], as well as strategies to maintain cooperation in the asymmetric prisoner's dilemma[41].

We can classify public goods games according to two properties, linearity and symmetry. We say a public goods game is linear if payoffs $u_i(\boldsymbol{e}, \boldsymbol{x})$ are linear in both arguments, $\boldsymbol{e}$ and $\boldsymbol{x}$. A public goods game is symmetric if players are indistinguishable, except for their endowments and for their contributions. Formally, if $\sigma = (\sigma_1, \ldots, \sigma_n)$ is a permutation of the numbers $1, \ldots, n$ and if $\boldsymbol{e}_\sigma$ and $\boldsymbol{x}_\sigma$ are the permuted endowment and contribution vectors, respectively, then the public goods game is symmetric if $u_i(\boldsymbol{e}, \boldsymbol{x}) = u_{\sigma_i}(\boldsymbol{e}_\sigma, \boldsymbol{x}_\sigma)$ for all permutations σ, endowments $\boldsymbol{e}$ and contributions $\boldsymbol{x}$. That is, if players were to switch roles with respect to their endowments and contributions, their payoffs would change accordingly. In particular, the public goods game in equation (1) is symmetric only if all players have the same productivity, $r_1 = \ldots = r_n$.

Equilibrium analysis. We explore under which conditions full cooperation can be sustained if the public goods game is repeated. After each round, there is another round with probability δ. Strategies for the repeated game are rules that tell the player which fraction of the endowment to contribute, depending on the players' endowments and on all previous contributions. If the contribution vector in round t is $\boldsymbol{x}(t)$, payoffs are given by the weighted average payoff per round, $\pi_i = (1 - \delta) \sum_{t=0}^{\infty} \delta^t u_i(\boldsymbol{e}, \boldsymbol{x}(t))$.

For a given public goods game with payoff function $u = (u_1, \ldots, u_n)$, continuation probability δ and endowment distribution $\boldsymbol{e}$, we say full cooperation is feasible if there is a subgame perfect equilibrium in which all players contribute their full endowment in every round. In a subgame perfect equilibrium, no player has an incentive to deviate after any given history[21]. It is a refinement of the Nash equilibrium concept: every subgame perfect equilibrium is a Nash equilibrium, but the converse does not need to be true.

We refer to the set of all endowments for which full cooperation is feasible as $E_u(\delta)$. In Fig. 2a–c, these sets are illustrated as blue areas within the space of all endowment distributions. In the Supplementary Information, we characterize these sets for all games that satisfy the four conditions (C), (PE), (IF) and (OC). We show that the following are equivalent (Supplementary Information, proposition 1): (i) cooperation is feasible for a given endowment distribution $\boldsymbol{e}$; (ii) the condition $\delta(u_i(\boldsymbol{e}, \boldsymbol{1}_{-i}) - u_i(\boldsymbol{e}, \boldsymbol{0})) \geq u_i(\boldsymbol{e}, \boldsymbol{1}_{-i}) - u_i(\boldsymbol{e}, \boldsymbol{1})$ holds for all players i with positive endowment. Here, $\boldsymbol{1}_{-i}$ is the shorthand notation for a group in which everyone contributes the full endowment, except for player i who

contributes nothing; and (iii) Grim is a subgame perfect equilibrium for the endowment distribution $\boldsymbol{e}$.

On the basis of this general characterization, we prove the following implications.

First, for any given payoff function u and continuation probability δ, there is a threshold $e_i^* < 1$ such that $\boldsymbol{e} \notin E_u(\delta)$ holds for any endowment distribution $\boldsymbol{e}$ with $e_i > e_i^*$. That is, cooperation is never feasible if one of the players receives an excessive share of the endowment.

Second, if the public goods game is symmetric and linear, and $E_u(\delta) \neq \varnothing$, then $(1/n, \ldots, 1/n) \in E_u(\delta)$. That is, if full cooperation is feasible in a linear and symmetric public goods game, then it is always feasible when all players receive the same endowment.

Third, if the public goods game is either asymmetric or nonlinear, there are cases for which $E_u(\delta) \neq \varnothing$, but $(1/n, \ldots, 1/n) \notin E_u(\delta)$. That is, in asymmetric or nonlinear public goods games, full cooperation may be feasible only for unequal endowment distributions.

Evolutionary analysis. We also explored the dynamics that arises if players have not yet settled on a particular equilibrium. Instead, they may begin with randomly initialized strategies, and then learn to use more profitable strategies over time.

To this end, we first consider a simplified strategy space. We assume that the contributions of the players in any given round depend only on the outcome of the previous round, as in most previous work on evolution of reciprocity[41–63]. In addition, we assume that players can choose only from among a fixed finite set $X = \{\hat{x}_1, \ldots, \hat{x}_m\}$ of possible contributions. Under these two assumptions, the strategies of the players take the form of a vector $\boldsymbol{p} = (p_{0,k}, p_{\boldsymbol{x},k})$. The entries $p_{0,k}$ for $1 \leq k \leq m$ give the probability of the player to choose contribution level $\hat{x}_k$ in the first round, when no previous history is yet available. The other entries $p_{\boldsymbol{x},k}$ give the probability of the player to choose $\hat{x}_k$ in subsequent rounds, conditional on the contribution vector $\boldsymbol{x} \in X^n$ of the previous round. For $\boldsymbol{p}$ to be a sensible strategy, we require $\sum_{k=1}^{m} p_{0,k} = 1$ and $\sum_{k=1}^{m} p_{\boldsymbol{x},k} = 1$ for all $\boldsymbol{x}$; that is, the strategy must prescribe an action for any given outcome of the previous round. When all players apply memory-1 strategies, their payoffs in the repeated game can be computed efficiently by representing the game as a Markov chain. The algorithm is shown in the Supplementary Information.

In the special case that players can give only their full endowment or nothing at all, we obtain $X = \{0, 1\}$. We refer to these two possible actions as 'cooperation' (C) and 'defection' (D). When there are only these two possible contribution levels, we can drop the index k in the definition of a memory-1 strategy and write $\boldsymbol{p} = (p_0; p_{\boldsymbol{x}})$. Under this notation, p_z is now the probability that the player will cooperate in the next round. If the game involves only two players, we obtain the typical format of memory-1 strategies for the iterated prisoner's dilemma[3], $\boldsymbol{p} = (p_0; p_{CC}, p_{CD}, p_{DC}, p_{DD})$. For example, the strategy Grim may be approximated by the memory-1 strategy $\boldsymbol{p} = (1; 1, 0, 0, 0)$. In the absence of errors, this memory-1 strategy cooperates only if both players have cooperated in all previous rounds.

Pure memory-1 strategies have entries that are either zero or one. Given the outcome of the previous round, their action is deterministic. Stochastic memory-1 strategies have entries that can take arbitrary values between zero and one. For given m and n, there are finitely many pure memory-1 strategies, but infinitely many stochastic memory-1 strategies.

For our evolutionary analysis, the actions of the players may be subject to implementation errors. That is, if the set of possible contributions is $X = \{\hat{x}_1, \ldots, \hat{x}_m\}$ and the player decides to choose the contribution $\hat{x}_k$, then the player will instead make a different contribution $\hat{x}_l$ with probability $\varepsilon/(m-1)$. We refer to ε as the error rate of the players. For infinitely repeated games (with $\delta = 1$), errors have the useful mathematical property that they make the game dynamics ergodic. As a result, the payoffs of the players will be independent of the contribution of the players in the very first round. In that case, we no longer need to specify a player's initial contribution probability $p_{0,k}$.

To model how players adapt their memory-1 strategies over time, we introduce an evolutionary process that we call 'introspection dynamics'. For this process, we again consider n players who interact in an asymmetric public goods game. In each evolutionary time step, one of the players is chosen at random to revise their strategy. To this end, this player i considers a randomly chosen alternative memory-1 strategy. Suppose the original strategy yields payoff π_i, whereas the alternative strategy yields payoff $\tilde{\pi}_i$ (keeping the strategies of the other players fixed). Then player i switches to the alternative strategy with probability $\rho = (1 + e^{-s(\tilde{\pi}_i - \pi_i)})^{-1}$. The parameter $s \geq 0$ represents the 'strength of selection'. In the limiting case $s \to 0$, the switching probability simplifies to $\rho = 1/2$, such that players adopt new strategies at random. In the other limiting case $s \to \infty$, the player adopts the alternative strategy only if it yields at least the payoff of their original strategy. Iterating these updates over many time steps, we obtain an ergodic process on the space of all strategy choices of n players. In particular, we note that this process has no absorbing states. Even if a strategy profile is an equilibrium, there is always a positive chance that one of the players deviates owing to chance. For small n, the invariant distribution of the evolutionary process can be calculated exactly. For larger n, the invariant distribution can be approximated by simulations.

Although memory-1 strategies have been routinely used to explore the evolution of reciprocity[41–63], it is natural to ask to what extent our results depend on the assumption of one-round memory. To explore this issue, we repeated all simulations with a more-general strategy space. We follow a previously published approach[15,64,65]. Players can choose among all strategies that can be represented by finite-state automata over the possible contributions X. Finite-state automata contain the previously considered memory-1 strategies as a special case. However, they can also encode strategies with arbitrarily long memory (Extended Data Fig. 6a). For the evolutionary process, we assume that when a mutation occurs, four cases can occur: (i) the action chosen in a given state changes, (ii) a transition between two states changes, (iii) a new state is added to the finite-state automaton or (iv) an existing state is removed (Extended Data Fig. 6b). Simulations show that although absolute cooperation rates tend to be lower for this strategy space, all of our qualitative predictions remain unchanged (Extended Data Fig. 6c). See Supplementary Information for details.

Experimental methods. For our experiment, we have recruited 436 participants on Amazon Mechanical Turk to take part in an interactive game. The experiment was implemented with SoPHIE, an online platform that allows for real-time interaction between participants on Amazon Mechanical Turk[10,66,67].

Participants were matched in pairs, which were randomly assigned to one of the five treatments. For each pair, one participant was randomly determined to adopt role A, whereas the other participant obtained role B. Players received US$1.00 for participating and could earn a bonus payment depending on their performance in the game. The tokens earned during the game were converted to US dollars at a rate of 800 tokens = US$1.00. The average bonus participants earned was US$1.70. After reading the experimental instructions (Supplementary Information), all participants had to pass a series of comprehension questions to ensure they understood the consequences of their decisions. All players were anonymous. They were identified only by their player ID (A or B). Each game consisted of at least 20 rounds. Thereafter, the game was continued with a 50% probability after each round to avoid end-game effects.

The behavioural experiment is based on the public goods game with the payoff function given by equation (1). Before the first round, both players were assigned an endowment e_i and a productivity value r_i. The possible values of e_i and r_i are depicted in Fig. 4a. Once assigned, the e_i and r_i of each participant remained constant throughout the experiment. Both players were informed about their own and the other player's endowment and productivity. Each round, participants decided how much to contribute to the public good. They could contribute any integer between 0 and e_i. A player's absolute contribution was multiplied by the respective productivity value r_i. All multiplied contributions were split equally among the players. Participants could not observe the other player's contribution before making their own decision. However, after each round, participants learned each other's contributions as well as the resulting payoffs for each player.

We analysed the data using two-tailed non-parametric tests, using pairs of two interacting players as our statistical unit. That is, for each quantity of interest, we calculated the respective average value for each pair of players; then we compared this average value across treatments (Fig. 4) or within each treatment (Extended Data Figs. 7, 10). For comparisons between treatments, we use the Mann–Whitney U-test, whereas for comparisons within a treatment we use the Wilcoxon signed-rank test. In the main text and all figures, we report the outcome of each test directly, without correcting for multiple testing. However, all of our key findings continue to hold when we apply Bonferroni correction (Supplementary Information section 5.3).

The sample size was determined in advance based on similar past research[10]. The number of groups that completed the experiment were $n_1 = 42$, $n_2 = 42$, $n_3 = 40$, $n_4 = 39$, $n_5 = 40$ for each of the five treatments, respectively. We find no significant differences between groups that completed the experiment and groups for which at least one player dropped out (Supplementary Information). For the statistical results presented in the main text, we used only the first 20 rounds of groups that completed the experiment. However, all of our conclusions remain valid if we include dropout groups by using multiple imputation (Supplementary Tables 1, 2).

In the Supplementary Information, we provide a full description of the methods used, and we report all test statistics and P values. Moreover, we discuss further aspects of our empirical results, such as the game dynamics over time or the distribution of contributions (Extended Data Figs. 7–10).

Parameters used for figures. Figure 2a–c shows the region in the endowment space in which full cooperation is feasible. The respective calculations are provided in the Supplementary Information. The first two columns are based on the linear public goods game with parameters $r_1 = r_2 = r_3 = 2$, $\delta = 0.8$ (Fig. 2a) and $r_1 = 2.9$, $r_2 = 1.5$, $r_3 = 1.1$, $\delta = 0.3$ (Fig. 2b), respectively. The last column considers a non-linear three-player public goods game with $\delta = 0.35$ (Fig. 2c) and payoff function

$$u_i(\boldsymbol{e}, \boldsymbol{x}) = \frac{1}{2}\max_{j,\,k \neq j}(e_j x_j + e_k x_k) + \frac{1}{3}\sum_{j=1}^{3} e_j x_j + (1-x_i)e_i \qquad (5)$$

This game represents a situation in which the two highest absolute contributions are of particular importance for the public good.

Figure 2d–f shows the outcome of evolutionary simulations. We systematically varied the players' initial endowments, considering all endowment distributions (e_1, e_2, e_3) for which $e_i \in \{0.00, 0.05, \ldots, 0.95, 1.00\}$. We used the same three payoff functions as in Fig. 2a–c, a continuation probability of $\delta = 1$ and strong selection, $s = 1,000$. Players use stochastic memory-1 strategies without errors, and they either contribute their full endowment or nothing, $X = \{0, 1\}$. The evolutionary process was simulated for 10^6 elementary time steps.

Figure 3 shows simulations as we vary player 1's endowment $e_1 \in \{0.00, 0.05, \ldots, 1.00\}$ and $e_2 = 1 - e_1$. We use the same productivity values as in the experiment, $r_1 = r_2 = 1.6$ (Fig. 3a) or $r_1 = 1.9$ and $r_2 = 1.3$ (Fig. 3b), respectively, and consider the case $\delta = 1$ and $s = 1,000$. To explore the robustness of our evolutionary findings, we consider three different scenarios, depending on the possible contribution levels in each round, $X_1 = \{0, 1\}$, $X_2 = \{0, 1/2, 1\}$ and $X_3 = \{0, 1/3, 2/3, 1\}$. Players can choose among all pure memory-1 strategies, subject to an error rate of $\varepsilon = 0.001$. For each value of e_1, simulations were run for at least 2×10^6 time steps for each data point.

Extended Data Figures 1–4 show further evolutionary results for the pairwise game considered in the behavioural experiment. For these figures, we assume that players choose only among pure memory-1 strategies with errors, and that they contribute only their full endowment or nothing in any given round. As a consequence, there are 16 possible strategies $\boldsymbol{p} = (p_{CC}, p_{CD}, p_{DC}, p_{DD})$. For example, the strategy of unconditional defectors is given by ALLD $= (0, 0, 0, 0)$, whereas WSLS[3] takes the form WSLS $= (1, 0, 0, 1)$. For these 16 strategies, we can compute numerically exact strategy abundances (Supplementary Information). Except for the parameters explicitly varied, all payoff parameters are the same as in the five experimental treatments, using a continuation probability $\delta = 1$, selection strength $s = 1,000$ and error rate $\varepsilon = 0.05$.

Extended Data Figure 5 considers a public goods game in which players have the same productivity, but they yield different benefits from the public goods. The payoff function is given by

$$u_i(\boldsymbol{e}, \boldsymbol{x}) = q_i r \sum_j x_j e_j + (1-x_i)e_i \qquad (6)$$

Here, r is the common productivity of the players and $q_i \in [0, 1]$ is player i's share of the public good. The game is a social dilemma only for certain values of q_i (see Supplementary Information for details). We show numerically exact results for pure memory-1 strategies with errors, and possible contribution levels $X = \{0, 1\}$. We use the parameters $r = 1.6$, $s = 1,000$ and $\varepsilon = 0.05$.

For Extended Data Fig. 6, we repeat the simulations shown in Fig. 3 for the case in which players can choose their strategies from the set of finite-state automata. We use the same baseline parameters as in Fig. 3. However, simulations are run for longer (5×10^6 time steps) and the error rate has been set to $\varepsilon = 0$ to allow for neutral invasions as previously described[15,64,65].

The results of our behavioural experiment are shown in Fig. 4 and Extended Data Figs. 7–10. As auxiliary information, we also provide error bars that indicate the respective 95% confidence intervals in Fig. 4 and Extended Data Figs. 7, 10. For the theoretical predictions, we used simulations for stochastic memory-1 strategies and possible contributions $x \in \{0, 1\}$. As parameters, we chose $s = 1,000$ and $\varepsilon = 0.001$. As indicated in Fig. 3 and Extended Data Fig. 4, our qualitative predictions are independent of the evolutionary parameters that we use, and independent of the possible contribution levels. A detailed description of the methods applied and of the depicted results is provided in the Supplementary Information.

Reporting summary. Further information on research design is available in the Nature Research Reporting Summary linked to this paper.

Data availability

The experimental data on which Fig. 4 and Extended Data Figs. 7–10 are based, as well as the STATA and R files that contain our statistical analysis, are available at https://osf.io/92jyw/.

Code availability

All evolutionary simulations and numerical calculations have been performed with MATLAB R2014A. We provide the respective scripts in the Supplementary Information. These scripts can be used to compute the payoffs of the players, to simulate the introspection dynamics and to numerically compute the expected dynamics.

31. Milinski, M., Sommerfeld, R. D., Krambeck, H.-J., Reed, F. A. & Marotzke, J. The collective-risk social dilemma and the prevention of simulated dangerous climate change. *Proc. Natl Acad. Sci. USA* **105**, 2291–2294 (2008).
32. Pacheco, J. M., Santos, F. C., Souza, M. O. & Skyrms, B. Evolutionary dynamics of collective action in *N*-person stag hunt dilemmas. *Proc. R. Soc. B* **276**, 315–321 (2009).
33. Jacquet, J. et al. Intra- and intergenerational discounting in the climate game. *Nat. Clim. Change* **3**, 1025–1028 (2013).
34. Vasconcelos, V. V., Santos, F. C. & Pacheco, J. M. A bottom-up institutional approach to cooperative governance of risky commons. *Nat. Clim. Change* **3**, 797–801 (2013).
35. Archetti, M. & Scheuring, I. Review: game theory of public goods in one-shot social dilemmas without assortment. *J. Theor. Biol.* **299**, 9–20 (2012).
36. Milinski, M., Röhl, T. & Marotzke, J. Cooperative interaction of rich and poor can be catalyzed by intermediate climate targets. *Clim. Change* **109**, 807–814 (2011).
37. Vasconcelos, V. V., Santos, F. C., Pacheco, J. M. & Levin, S. A. Climate policies under wealth inequality. *Proc. Natl Acad. Sci. USA* **111**, 2212–2216 (2014).
38. Abou Chakra, M. & Traulsen, A. Under high stakes and uncertainty the rich should lend the poor a helping hand. *J. Theor. Biol.* **341**, 123–130 (2014).
39. Abou Chakra, M., Bumann, S., Schenk, H., Oschlies, A. & Traulsen, A. Immediate action is the best strategy when facing uncertain climate change. *Nat. Commun.* **9**, 2566 (2018).
40. Hauser, O. P., Traulsen, A. & Nowak, M. A. Heterogeneity in background fitness acts as a suppressor of selection. *J. Theor. Biol.* **343**, 178–185 (2014).
41. Akin, E. What you gotta know to play good in the Iterated Prisoner's Dilemma. *Games* **6**, 175–190 (2015).
42. Nowak, M. A. & Sigmund, K. Tit for tat in heterogeneous populations. *Nature* **355**, 250–253 (1992).
43. Frean, M. R. The prisoner's dilemma without synchrony. *Proc. R. Soc. Lond. B* **257**, 75–79 (1994).
44. Killingback, T., Doebeli, M. & Knowlton, N. Variable investment, the Continuous Prisoner's Dilemma, and the origin of cooperation. *Proc. R. Soc. Lond. B* **266**, 1723–1728 (1999).
45. Imhof, L. A. & Nowak, M. A. Stochastic evolutionary dynamics of direct reciprocity. *Proc. R. Soc. B* **277**, 463–468 (2010).
46. Kurokawa, S., Wakano, J. Y. & Ihara, Y. Generous cooperators can outperform non-generous cooperators when replacing a population of defectors. *Theor. Popul. Biol.* **77**, 257–262 (2010).
47. García, J. & Traulsen, A. The structure of mutations and the evolution of cooperation. *PLoS ONE* **7**, e35287 (2012).
48. Grujić, J., Cuesta, J. A. & Sánchez, A. On the coexistence of cooperators, defectors and conditional cooperators in the multiplayer iterated Prisoner's Dilemma. *J. Theor. Biol.* **300**, 299–308 (2012).
49. Press, W. H. & Dyson, F. J. Iterated Prisoner's Dilemma contains strategies that dominate any evolutionary opponent. *Proc. Natl Acad. Sci. USA* **109**, 10409–10413 (2012).
50. Van Segbroeck, S., Pacheco, J. M., Lenaerts, T. & Santos, F. C. Emergence of fairness in repeated group interactions. *Phys. Rev. Lett.* **108**, 158104 (2012).
51. Akin, E. in *Ergodic Theory, Advances in Dynamics* (ed. Assani, I.) 77–107 (de Gruyter, Berlin, 2016).
52. Stewart, A. J. & Plotkin, J. B. Collapse of cooperation in evolving games. *Proc. Natl Acad. Sci. USA* **111**, 17558–17563 (2014).
53. Stewart, A. J. & Plotkin, J. B. The evolvability of cooperation under local and non-local mutations. *Games* **6**, 231–250 (2015).
54. Szolnoki, A. & Perc, M. Defection and extortion as unexpected catalysts of unconditional cooperation in structured populations. *Sci. Rep.* **4**, 5496 (2014).
55. Toupo, D. F. P., Rand, D. G. & Strogatz, S. H. Limit cycles sparked by mutation in the repeated prisoner's dilemma. *Int. J. Bifurc. Chaos* **24**, 1430035 (2014).
56. Dong, Y., Li, C., Tao, Y. & Zhang, B. Evolution of conformity in social dilemmas. *PLoS ONE* **10**, e0137435 (2015).
57. Pan, L., Hao, D., Rong, Z. & Zhou, T. Zero-determinant strategies in iterated public goods game. *Sci. Rep.* **5**, 13096 (2015).
58. Baek, S. K., Jeong, H. C., Hilbe, C. & Nowak, M. A. Comparing reactive and memory-one strategies of direct reciprocity. *Sci. Rep.* **6**, 25676 (2016).
59. McAvoy, A. & Hauert, C. Autocratic strategies for iterated games with arbitrary action spaces. *Proc. Natl Acad. Sci. USA* **113**, 3573–3578 (2016).
60. Reiter, J. G., Hilbe, C., Rand, D. G., Chatterjee, K. & Nowak, M. A. Crosstalk in concurrent repeated games impedes direct reciprocity and requires stronger levels of forgiveness. *Nat. Commun.* **9**, 555 (2018).
61. Ichinose, G. & Masuda, N. Zero-determinant strategies in finitely repeated games. *J. Theor. Biol.* **438**, 61–77 (2018).
62. Hilbe, C., Chatterjee, K. & Nowak, M. A. Partners and rivals in direct reciprocity. *Nat. Hum. Behav.* **2**, 469–477 (2018).
63. Hilbe, C., Šimsa, Š., Chatterjee, K. & Nowak, M. A. Evolution of cooperation in stochastic games. *Nature* **559**, 246–249 (2018).
64. García, J. & van Veelen, M. In and out of equilibrium I: evolution of strategies in repeated games with discounting. *J. Econ. Theory* **161**, 161–189 (2016).
65. García, J. & van Veelen, M. No strategy can win in the repeated Prisoner's Dilemma: linking game theory and computer simulations. *Front. Robot. AI* **5**, 102 (2018).
66. Hendriks, A. *SoPHIE — Software Platform for Human Interaction Experiments*. https://www.sophie.uni-osnabrueck.de/start/ (2012).
67. Hauser, O. P., Hendriks, A., Rand, D. G. & Nowak, M. A. Think global, act local: preserving the global commons. *Sci. Rep.* **6**, 36079 (2016).

Author contributions All authors conceived the study, performed the analysis, discussed the results and wrote the manuscript.

**Competing interests** The authors declare no competing interests.

Additional information
Supplementary information is available for this paper at https://doi.org/10.1038/s41586-019-1488-5.
Correspondence and requests for materials should be addressed to O.P.H. or C.H. or M.A.N.
Peer review information *Nature* thanks Joe Yuichiro Wakano and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.
**Reprints and permissions information** is available at http://www.nature.com/reprints.
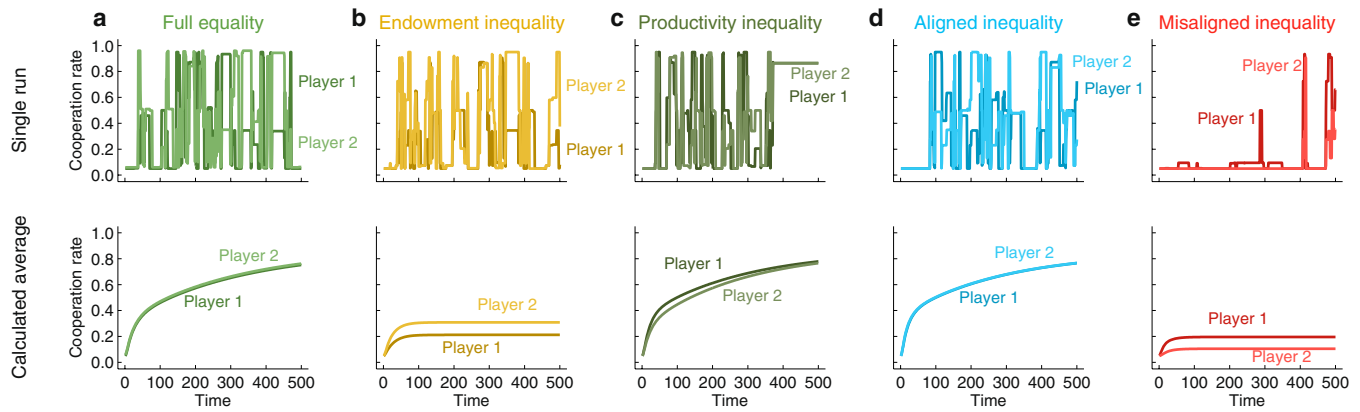
Extended Data Fig. 1 | Dynamics of cooperation among unequals.
In Figs. 2, 3, we show how often players cooperate on average. **a–e,** Here
we depict evolutionary trajectories over time, for the five treatments
considered in our experiment. We assume that players can choose among
the 16 pure memory-1 strategies. Top, five single runs of the introspection
dynamics. Bottom, expected trajectories of the introspection dynamics,
which can be derived explicitly (Supplementary Information section 4.3).
These expected trajectories represent the cooperation rate over time as
we average over many realizations of the process. We observe substantial
cooperation in three of the five cases: in the treatments with full equality
(**a**), productivity inequality (**c**) and aligned inequality (**d**).
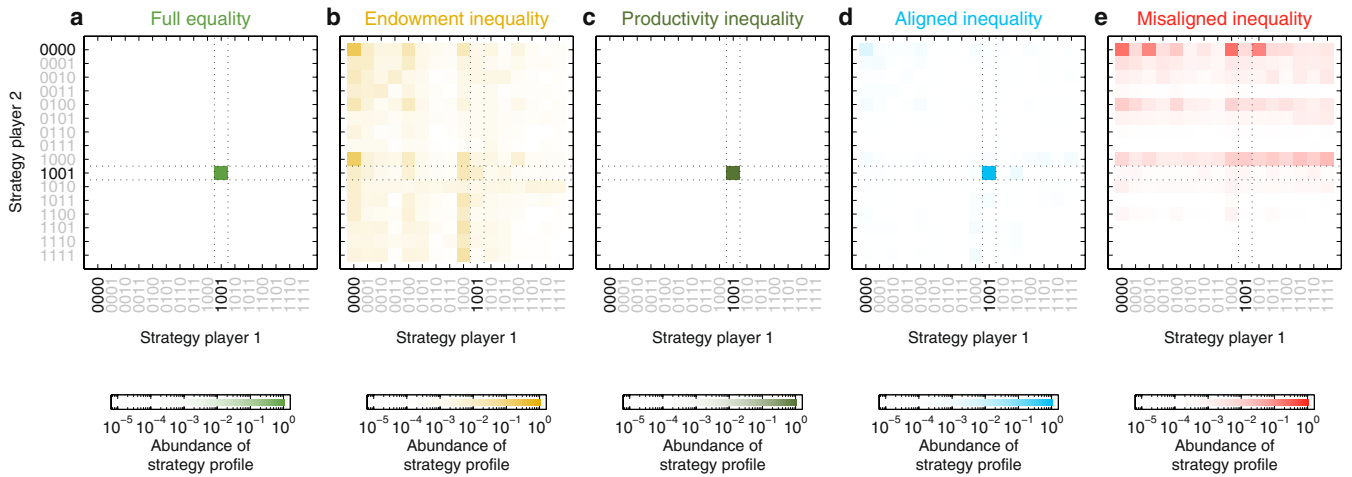
**Extended Data Fig. 2 | Under endowment inequality and misaligned inequality, players fail to coordinate on WSLS.** Here, we consider the long-run dynamics of the games considered in Extended Data Fig. 1. For each pair (p_1, p_2) of pure memory-1 strategies, we can compute how often the respective strategy pair is played according to the invariant distribution of the evolutionary process. **a**, **c**, **d**, Under full equality, productivity inequality or aligned inequality, players typically coordinate on a WSLS equilibrium, as indicated by the coloured square in the centre of the dotted lines. **b**, **e**, Under endowment inequality or misaligned inequality, players fail to coordinate on a unique equilibrium. Instead, most of the evolving strategies prescribe to defect against the opponent. We note that in those treatments in which players have different endowments, the low-endowment player faces a reduced strength of selection (because the endowment of this player is reduced from 0.5 to 0.25). As a consequence, the marginal distribution of the low-endowment player in **b**, **e** is more uniform than the marginal distribution of the high-endowment player.
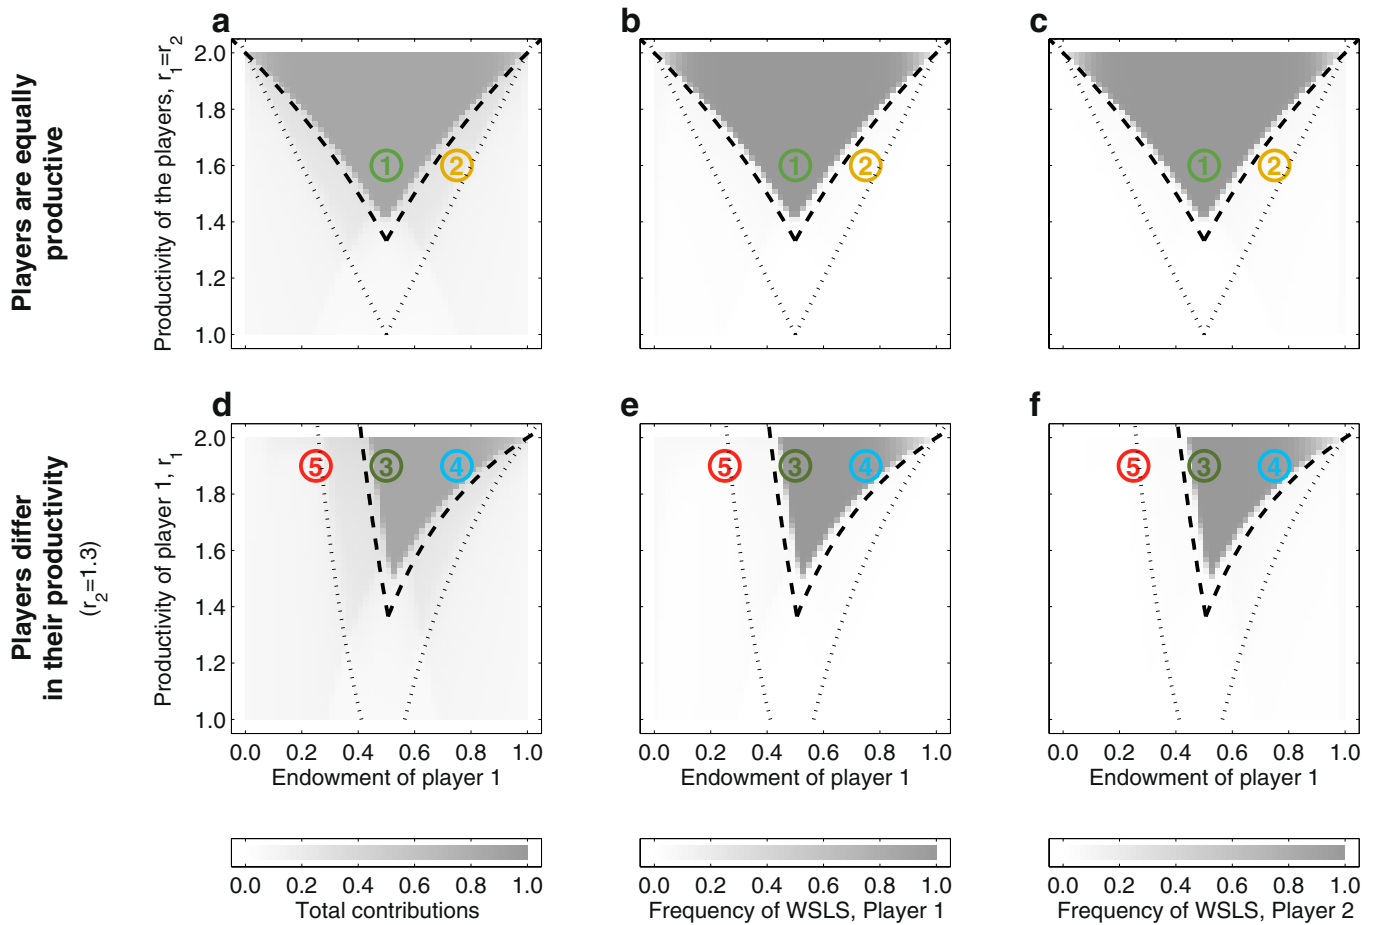
Extended Data Fig. 3 | An equilibrium analysis explains why cooperation emerges in only three of the five treatments. Using the same two-player setup as in Extended Data Figs. 1, 2, we explored how much players contribute on average when we simultaneously vary the endowment (x axis) as well as their productivity r_1 of player 1. For each parameter combination, we record the total contributions of the player and how often they use WSLS according to the invariant distribution of the evolutionary process (indicated in shades of grey). We compare these evolutionary results with the region for which WSLS is an equilibrium (indicated by dashed lines) and with the region for which Grim is an equilibrium (dotted lines); see Supplementary Information for details. The coloured symbols indicate which parameter combinations have been used for the experimental treatments. **a–c**, For equal productivities, the full equality treatment (1) is in the region in which cooperation can evolve, whereas the unequal endowment treatment (2) is not. **d–f**, For unequal productivities, only the misaligned inequality treatment (5) is outside the region in which cooperation can evolve.

Full equality · Endowment inequality · Productivity inequality · Aligned inequality · Misaligned inequality
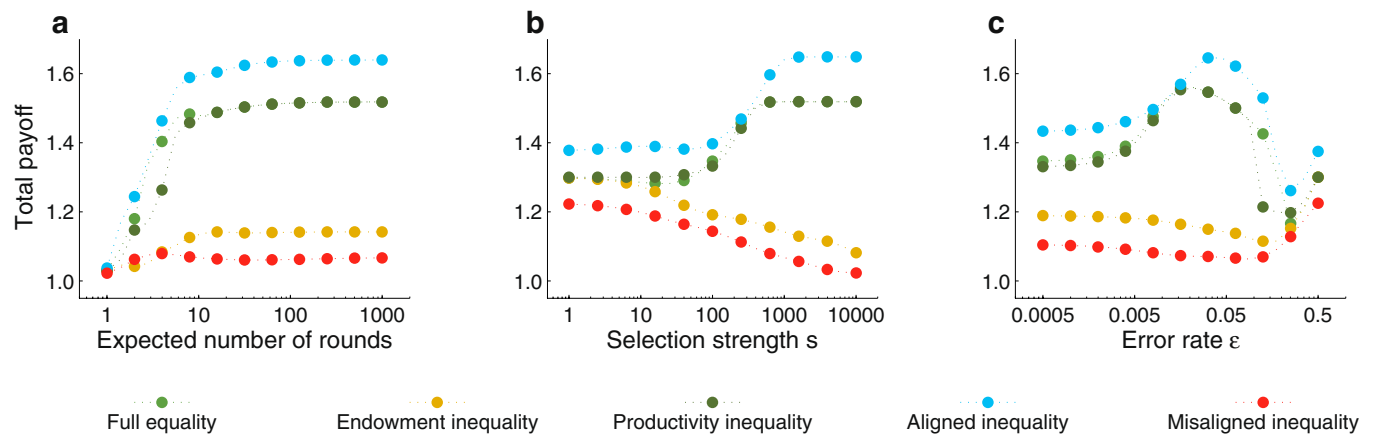
Extended Data Fig. 4 | Robustness of evolutionary results with respect to parameter changes. a–c, To explore the robustness of our theoretical predictions, we varied the expected number of rounds played between two players (**a**), the selection strength (**b**) and the rate at which players commit an implementation error (**c**). Although the quantitative results depend on these parameters, the qualitative ordering of the five treatments is the same across all considered scenarios. Except for the parameters explicitly varied on the x axis, all parameters are the same as in Extended Data Figs. 1, 2.
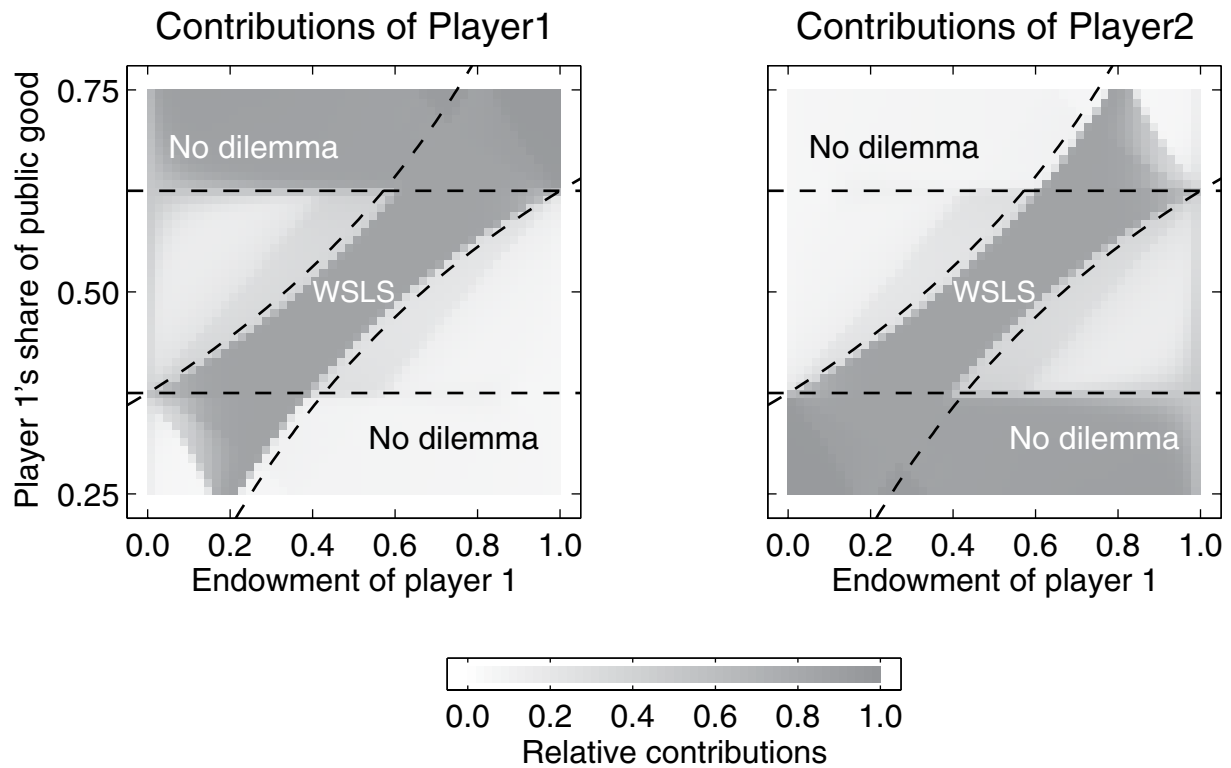
Extended Data Fig. 5 | Cooperation in an asymmetric game in which players derive different payoffs from the public good. Instead of considering players who differ in their productivity, here we consider an asymmetric two-player public goods game in which players differ in the share of the public goods that they get (the exact model is specified in the Supplementary Information). We vary two parameters, player 1's share of the initial endowment, and player 1's share of the public good. For each parameter combination, we record the average contributions of the players over the course of the evolutionary process (indicated in the grey colour). For games in which players get different shares of the public good, we note that the game is a social dilemma only if neither player's share is too large (otherwise that player would always have an incentive to cooperate, no matter what the co-player does). However, if both players get an intermediate share of the public good, full cooperation can again evolve when WSLS is an equilibrium.
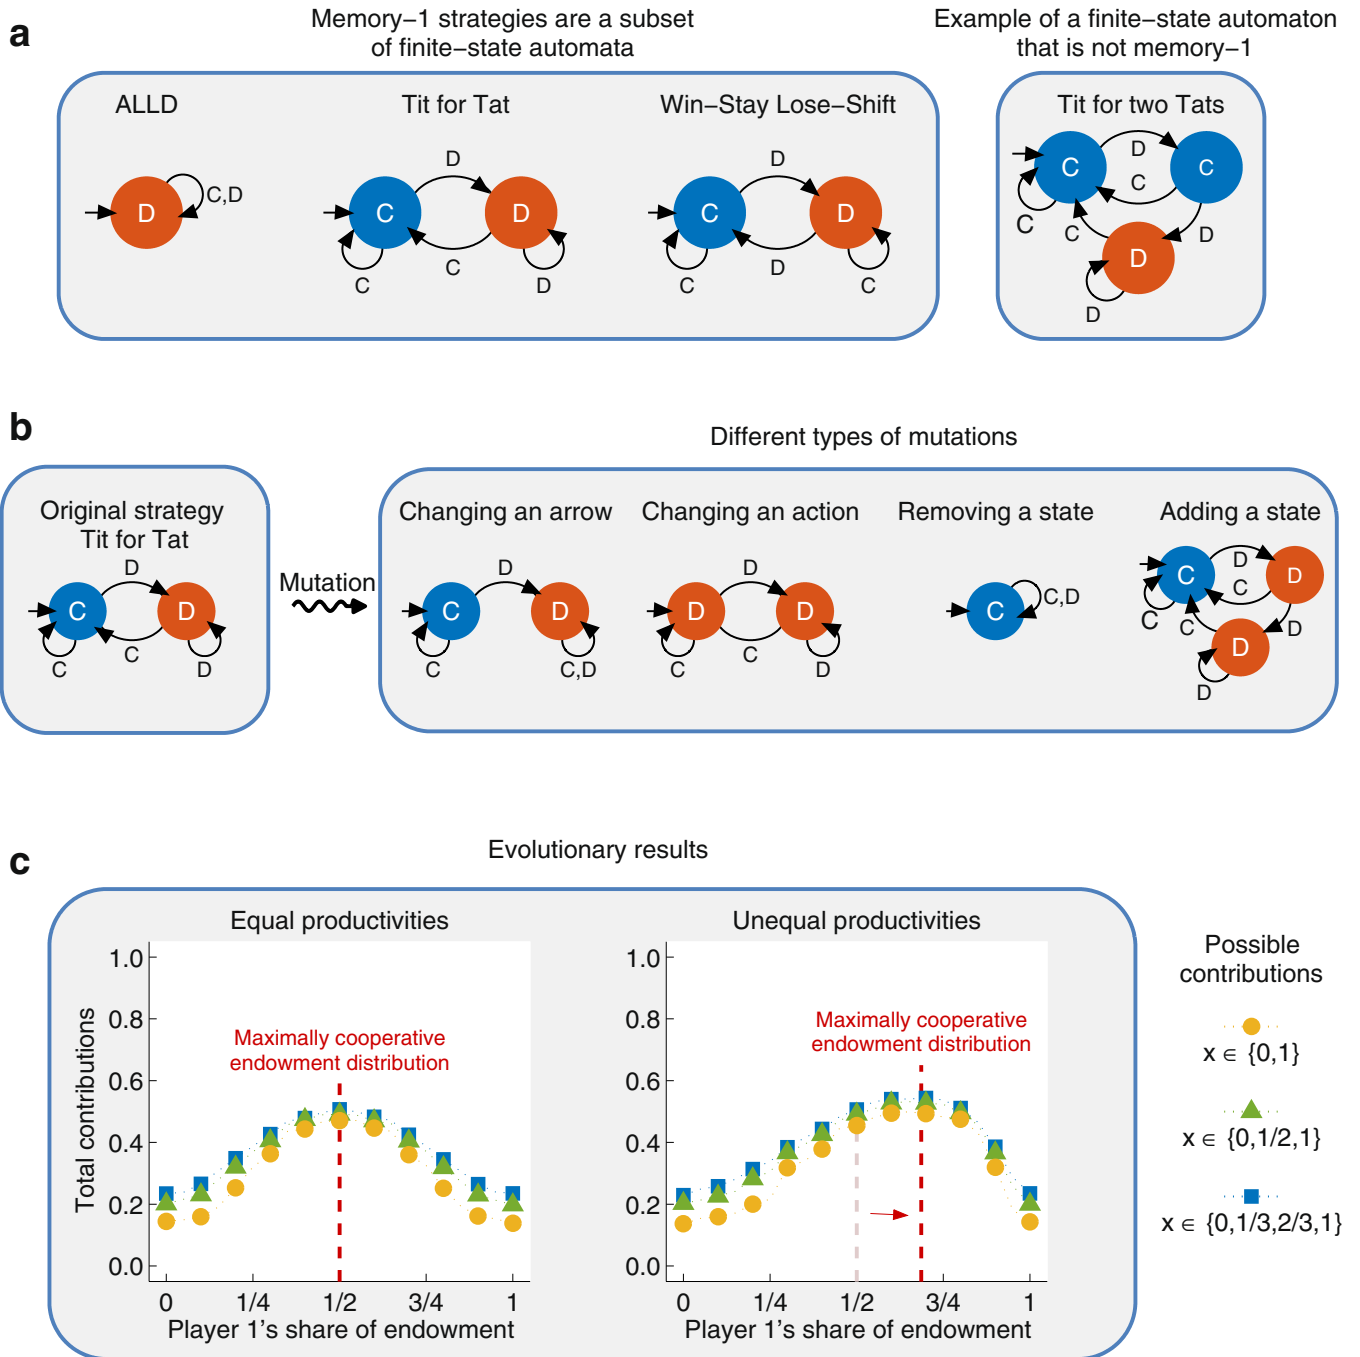
Extended Data Fig. 6 | Evolution of cooperation among players using finite-state automata. a, Here we represent finite-state automata for a game between two players in which players can either contribute their full endowment (C) or nothing (D). A finite-state automaton consists of three components: a set of states (represented by the large circles), the action played in each state (represented by the colour of the circle and the letters 'C' and 'D') and a transition rule (represented by arrows; the associated letter shows for which of the co-player's actions the respective arrow is taken). Finite-state automata are able to implement all memory-1 strategies. In addition, they can encode strategies that depend on arbitrarily long sequences of past actions. **b**, To model evolution among finite-state automata, we use a previously published mutation scheme[15,64]. When a mutation occurs, the direction of a random arrow is changed, the action in a randomly chosen state is changed, a random state is removed or a state is added. **c**, Using this more general strategy space, we repeated the simulations in Fig. 3. Although overall cooperation rates are slightly lower, all qualitative results remain unchanged.
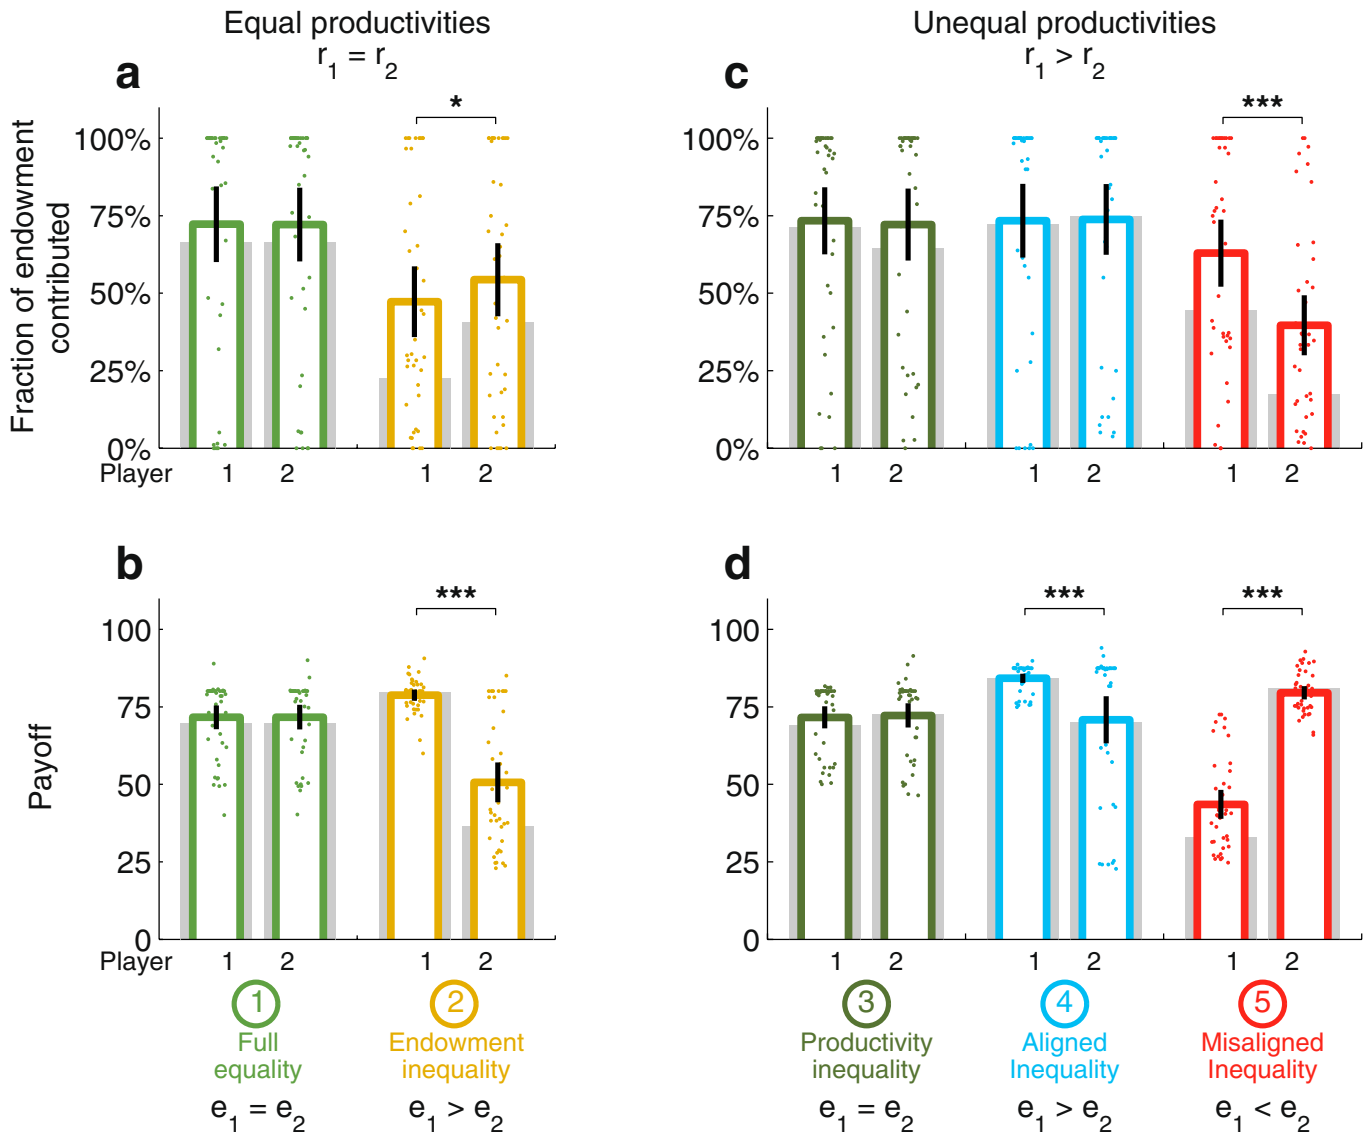
Extended Data Fig. 7 | Contributions and payoffs of the two players across treatments. For each of the five experimental treatments, we compare the average contributions and the average payoff of the two players. Grey bars indicate the theoretical prediction based on evolutionary simulations. Coloured bars depict the outcome of the experiment. Error bars represent the respective 95% confidence intervals. Asterisks indicate statistical differences based on two-tailed Wilcoxon signed-rank tests. The number of groups per treatment is 42, 42, 40, 39, 40 for treatments 1–5, respectively. **a, b**, Under full equality, the two players contribute a similar share of their endowment and they obtain approximately equal payoffs. Under endowment inequality, the cooperation rates of both players are reduced, with the contributions of the high-endowment player (player 1) being significantly lower than the contributions of player 2. **c, d**, For productivity inequality and aligned inequality, we find no differences in the relative contributions of the players. For misaligned inequality, the relative contributions of the better-endowed but less-productive player 2 are considerably reduced. For both aligned and misaligned inequality, the two players earn significantly different payoffs. Nevertheless, the player with the lower payoff in the aligned inequality treatment derives a similar payoff as the two player types under productivity inequality. For details, see Supplementary Information.
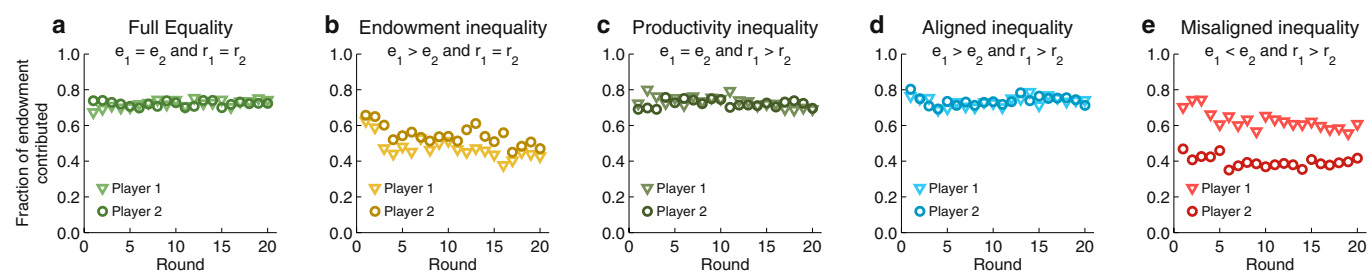
Extended Data Fig. 8 | Experimental dynamics of cooperation.
a–e, For each of the five treatments, we show the average contributions of the players over the course of the experiment. In all treatments the contributions are relatively stable over time, except for a significant negative trend in the treatment with endowment inequality (**b**) (see Supplementary Information for details).
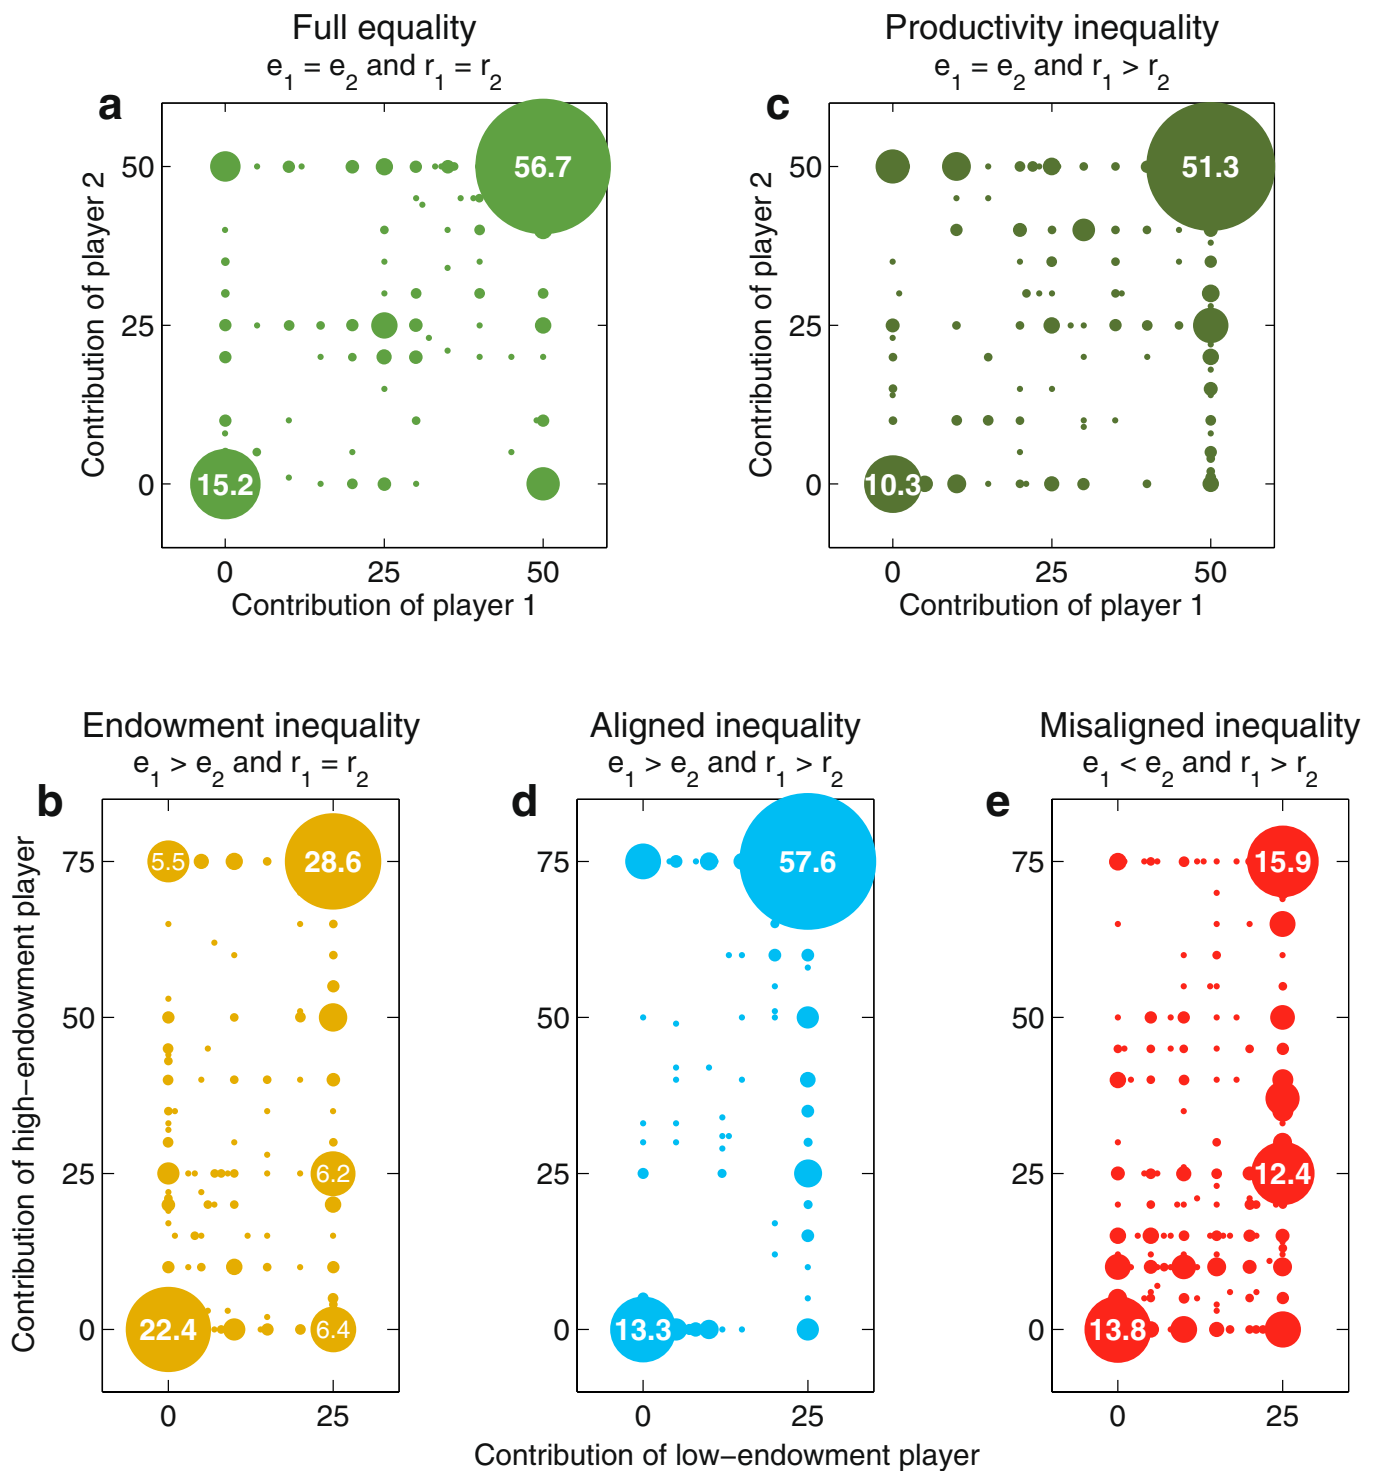
Extended Data Fig. 9 | Individual cooperation decisions across the five treatments. a–e, To analyse the joint contribution decisions of the two players, we plot here how often player 1 has contributed y_1 tokens while player 2 has contributed y_2 tokens, for each pair (y_1, y_2). **a, c, d,** Under full equality (**a**), productivity inequality (**c**) and aligned inequality (**d**), most individual decisions are mutually cooperative. **b, e,** By contrast, under endowment inequality (**b**) and misaligned inequality (**e**), contributions are more scattered. **e,** Moreover, in the treatment with misaligned inequality, we observe that a substantial fraction of high-endowment players only matches the absolute contributions of the other player. For example, in 12.4% of the rounds, the low-endowment player contributes all 25 tokens at their disposal, and the high-endowment player contributes the same absolute amount of tokens (corresponding to 1/3 of this player's endowment).

a Equal productivities $r_1 = r_2$

b Unequal productivities $r_1 > r_2$
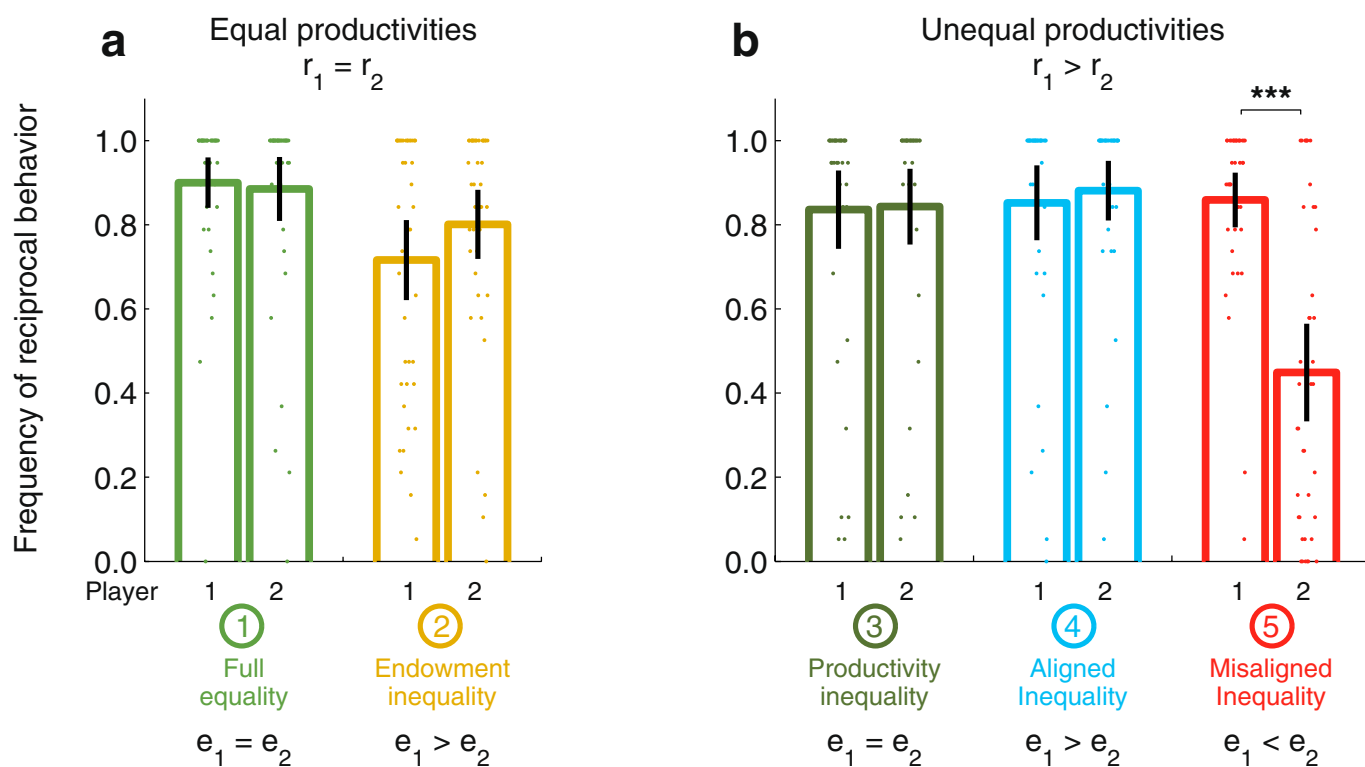
Extended Data Fig. 10 | Abundance of reciprocal behaviours across the five treatments. a, b, To explore whether subjects apply reciprocal strategies, we show the fraction of rounds in which subjects match or exceed their co-player's relative contribution from the previous round. That is, if player 1 has contributed x% of their endowment in round t, we record whether or not player 2 contributes at least x% of their endowment in round $t + 1$. Note that reciprocal strategies do not automatically yield high cooperation rates, because mutually defecting players are also reciprocal. Error bars represent the respective confidence intervals. Statistically significant differences were analysed using a two-tailed Wilcoxon signed-rank test. ***$P < 0.001$. Sample sizes are 42, 42, 40, 39, 40 for the treatments 1–5, respectively. Generally, we find high levels of reciprocity; only in the treatment with misaligned inequality does the high-endowment low-productivity player 2 exhibit a strongly reduced reciprocity rate. See Supplementary Information for details.

nature research

Corresponding author(s):   Martin A. Nowak

Reporting Summary

Nature Research wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Research policies, see Authors & Referees and the Editorial Policy Checklist.

Statistical parameters

When statistical analyses are reported, confirm that the following items are present in the relevant location (e.g. figure legend, table legend, main text, or Methods section).

n/a	Confirmed	
☐	☒	The exact sample size (n) for each experimental group/condition, given as a discrete number and unit of measurement
☐	☒	An indication of whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
☐	☒	The statistical test(s) used AND whether they are one- or two-sided *Only common tests should be described solely by name; describe more complex techniques in the Methods section.*
☒	☐	A description of all covariates tested
☐	☒	A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
☐	☒	A full description of the statistics including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
☐	☒	For null hypothesis testing, the test statistic (e.g. F, t, r) with confidence intervals, effect sizes, degrees of freedom and P value noted *Give P values as exact values whenever suitable.*
☒	☐	For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
☐	☒	For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
☐	☒	Estimates of effect sizes (e.g. Cohen's d, Pearson's r), indicating how they were calculated
☐	☒	Clearly defined error bars *State explicitly what error bars represent (e.g. SD, SE, CI)*

Our web collection on statistics for biologists may be useful.

Software and code

Policy information about availability of computer code

Data collection	All evolutionary simulations were performed using Matlab R2014A. The online experiment was implemented with SoPHIE, an online platform that allows for real-time interaction between Amazon Mechanical Turk (AMT) participants.
Data analysis	To analyze the experimental data, we used Stata SE 15.1, R 1.1.453, and Microsoft Excel 14.7.7.

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors/reviewers upon request. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Research guidelines for submitting code & software for further information.

Data

Policy information about availability of data

All manuscripts must include a data availability statement. This statement should provide the following information, where applicable:
- Accession codes, unique identifiers, or web links for publicly available datasets
- A list of figures that have associated raw data
- A description of any restrictions on data availability

The data is publicly available on OSF: https://osf.io/92jyw/

Field-specific reporting

Please select the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

☐ Life sciences ☐ Behavioural & social sciences ☒ Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see nature.com/authors/policies/ReportingSummary-flat.pdf

Ecological, evolutionary & environmental sciences study design

All studies must disclose on these points even when the disclosure is negative.

Study description	Participants were matched in pairs, which were assigned to one of five conditions (as described in Fig. 4a). The conditions varied across two dimensions: equality or inequality in endowments and equality or inequality in productivities. The five conditions are (1) Full equality in both dimensions, (2) Endowment inequality only, (3) Productivity inequality only, (4) Aligned inequality (the more productive player has a higher endowment), (5) Misaligned inequality (the more productive player has a lower endowment). Each pair of players interacted over at least 20 rounds. After that, the game was repeated with 50% probability after each round. See SI Section 5.1 for details. Because of the repeated interaction between two players, the data is structured hierarchically. For the main analysis, we averaged the two players' contributions at the round level and subsequently at the group level. The experiment was conducted exactly once. The number of participants for each condition is described below.
Research sample	We recruited N = 436 participants on Amazon Mechanical Turk (AMT).
Sampling strategy	The sample size was determined in advance based on similar past research (e.g. Hauser et al, Behavioral Public Policy, 2019). In total, we recruited 88, 88, 86, 86, and 88 subjects to participate in the five conditions, respectively. The number of subjects that completed the experiment was 84, 84, 80, 78, and 80, respectively. See SI Section 5 for details. The number of groups sampled has not been based on a formal power analysis. However, we chose our sample size such that we have a comparably large number of groups per condition, relative to previous studies on direct reciprocity (e.g. Wedekind and Milinski, PNAS 1996; Milinski and Wedekind, PNAS 1998; Dal Bo and Frechette, Am Econ Rev 2011; Fudenberg et al, Am Econ Rev 2012; Hilbe et al, Nature Communications 2015).
Data collection	The data was collected using SoPHIE, an online experimental software. Once the experiment is programmed and launched, data collection proceeds automatically. Once the data collection is completed, we downloaded the data for analysis offline.
Timing and spatial scale	Our experiments were conducted across 10 experimental sessions on AMT in October 2015.
Data exclusions	As planned from the outset, we only statistically analyzed the first twenty rounds of each game, because this was the maximum number of rounds all groups had in common. Moreover, we excluded groups in which at least one participant dropped out half-way through the game. However, all our results are unchanged when we use imputed values for dropout groups. See SI Section 5.3 for details.
Reproducibility	The experiment was conducted only once; however, our sample size is sufficiently large (and perhaps larger then usual) for experimental games of this nature.
Randomization	Randomization occurred at the session level.
Blinding	Blinding was not necessary because the experimental software took care of randomization and data collection automatically.

Did the study involve field work? ☐ Yes ☒ No

Reporting for specific materials, systems and methods

Materials & experimental systems

n/a	Involved in the study
☒	☐ Unique biological materials
☒	☐ Antibodies
☒	☐ Eukaryotic cell lines
☒	☐ Palaeontology
☒	☐ Animals and other organisms
☐	☒ Human research participants

Methods

n/a	Involved in the study
☒	☐ ChIP-seq
☒	☐ Flow cytometry
☒	☐ MRI-based neuroimaging

Human research participants

Population characteristics	Our study population is drawn from the general Amazon Mechanical Turk (see above) pool, which has been described in detail in other research (e.g. see Buhrmester et al. 2011)
Recruitment	Participants were recruited on AMT through a standard procedure by describing the nature of this research, the length of the task, the payoff for participating, and the potential for an additional bonus payment depending on decisions made during the study. For details, see SI Section 5.1.

Social dilemmas among unequals

Oliver P. Hauser[1,6]*, Christian Hilbe[2,6]*, Krishnendu Chatterjee[2] & Martin A. Nowak[3,4,5]*

[1]Department of Economics, University of Exeter Business School, Exeter, UK. [2]IST Austria, Klosterneuburg, Austria. [3]Program for Evolutionary Dynamics, Harvard University, Cambridge, MA, USA. [4]Department of Organismic and Evolutionary Biology, Harvard University, Cambridge, MA, USA. [5]Department of Mathematics, Harvard University, Cambridge, MA, USA. [6]These authors contributed equally: Oliver P. Hauser, Christian Hilbe. *e-mail: o.hauser@exeter.ac.uk; christian.hilbe@ist.ac.at; martin_nowak@harvard.edu

Supplementary Information
Social dilemmas among unequals

Oliver P. Hauser[1,*], Christian Hilbe[2,*], Krishnendu Chatterjee[2], Martin A. Nowak[3,4,5]

[1]Department of Economics, University of Exeter Business School, Exeter EX4 4PU, UK; [2]IST Austria, 3400 Klosterneuburg, Austria; [3]Program for Evolutionary Dynamics, [4]Department of Organismic and Evolutionary Biology, [5]Department of Mathematics, Harvard University, Cambridge MA 02138, USA.

* These authors contributed equally to this work.

Supplementary Methods and Results

In Section 1, we describe the basic setup of our model. We introduce a general framework to analyze repeated public good games among unequals, and we describe the players' possible strategies and the resulting payoffs. In Section 2, we summarize previous work on asymmetric social dilemmas, and we discuss the novelty of our approach. In Section 3 we analyze under which conditions players can sustain cooperation, depending on the players' endowments and on the shape of the payoff function. We complement this static equilibrium analysis with an evolutionary perspective in Section 4. There, we consider players who adapt their strategies over time. We explore under which conditions players learn to be fully cooperative, assuming that the initial population consists of unconditional defectors. In Section 5 we give a full description of our experimental study, including a more detailed analysis of the subjects' behavior. In Section 6, we critically review some of our model assumptions, and we propose several model extensions for future work. Finally, the appendix contains the proofs of our equilibrium results, the MATLAB code for the simulations, and the game instructions of our behavioral experiment.

1 Model description

1.1 A general framework for public good games among unequals

In the following, we introduce a general framework for social dilemmas with multiple players who may differ along multiple dimensions. To this end, let us consider a group of $n \geq 2$ players. Initially, each group member i is given an endowment $e_i \geq 0$. We refer to the n-tuple $\mathbf{e} = (e_1, \ldots, e_n)$ as the group's endowment vector. These endowments may be interpreted as the players' regular income, or as the amount of time they have at their free disposal. Without loss of generality, we assume total endowments are normalized such that $e_1 + \ldots + e_n = 1$. When players have equal endowments, the endowment vector is $\mathbf{e} = (1/n, \ldots, 1/n)$.

Each player can choose which fraction x_i of her endowment she wishes to contribute to the public good, with $0 \leq x_i \leq 1$. We call the n-tuple $\mathbf{x} = (x_1, \ldots, x_n)$ the group's contribution vector. Given an endowment vector $\mathbf{e}$ and a contribution vector $\mathbf{x}$, player i's absolute contribution to the public good is $e_i \cdot x_i$. We use the shortcut notation $\mathbf{x} = \mathbf{1}$ for the contribution vector where each player contributes her full endowment. We write $\mathbf{x} = \mathbf{0}$ when each player contributes nothing at all. For two contribution vectors $\mathbf{x} = (x_1, \ldots, x_n)$ and $\mathbf{x}' = (x_1', \ldots, x_n')$, we write $\mathbf{x} \geq \mathbf{x}'$ if $x_i \geq x_i'$ for all players i.

The payoffs of the public good game may depend on the endowments, the contributions, and on personal characteristics of the players (such as a player's productivity, or the degree to which a player benefits from publicly provided goods). Specifically, given the distribution of endowments $\mathbf{e}$, and the relative contributions $\mathbf{x}$, players yield the payoff $u(\mathbf{e}, \mathbf{x}) \in \mathbb{R}^n$, where u denotes the public good's payoff function. The i-th component $u_i(\mathbf{e}, \mathbf{x})$ corresponds to the payoff of player i. We assume that when no player contributes, players end up with their initial endowments, $u_i(\mathbf{e}, \mathbf{0}) = e_i$ for all players i. We refer to the sum over all individual payoffs, $U(\mathbf{e}, \mathbf{x}) := \sum_{i=1}^{n} u_i(\mathbf{e}, \mathbf{x})$, as the group payoff.

In the following, we will mostly be interested in situations in which there is a tension between individual and collective incentives to contribute. To formalize this tension, we consider public good games with the following four properties.

(C) Continuity: *Payoffs change continuously in the players' endowments and contributions.*

The payoff function $u(\mathbf{e}, \mathbf{x})$ is continuous in both its arguments, $\mathbf{e}$ and $\mathbf{x}$.

(PE) Positive externalities: *Players prefer their co-players to make higher contributions.*

If $\mathbf{x}$ and $\mathbf{x}'$ are two contribution vectors such that $x_i = x_i'$ and $x_j \geq x_j'$ for all $j \neq i$ then $u_i(\mathbf{e}, \mathbf{x}) \geq u_i(\mathbf{e}, \mathbf{x}')$ for all endowment distributions $\mathbf{e}$. In this case, the strict inequality $u_i(\mathbf{e}, \mathbf{x}) > u_i(\mathbf{e}, \mathbf{x}')$ holds if and only if there is at least one co-player j with $e_j > 0$ and $x_j > x_j'$.

(IF) Incentives to free-ride: *Players have an incentive to lower their own contributions.*

If $\mathbf{x}$ and $\mathbf{x}'$ are two contribution vectors such that $x_i < x_i'$ and $x_j = x_j'$ for all $j \neq i$ then $u_i(\mathbf{e}, \mathbf{x}) \geq u_i(\mathbf{e}, \mathbf{x}')$ for all endowment distributions $\mathbf{e}$. In this case, the strict inequality $u_i(\mathbf{e}, \mathbf{x}) > u_i(\mathbf{e}, \mathbf{x}')$ holds if and only if player i's endowment is positive, $e_i > 0$.

(OC) Optimality of cooperation: *From a collective perspective, higher contributions are preferred.*

If $\mathbf{x}$ and $\mathbf{x}'$ are two contribution vectors such that $\mathbf{x} \geq \mathbf{x}'$ then $U(\mathbf{e}, \mathbf{x}) \geq U(\mathbf{e}, \mathbf{x}')$. In this case, the strict inequality $U(\mathbf{e}, \mathbf{x}) > U(\mathbf{e}, \mathbf{x}')$ holds if and only if there is a player i with $e_i > 0$ and $x_i > x_i'$.

The first condition of continuity is mainly required for technical reasons; it will be useful to show that full cooperation cannot be sustained when initial endowments are too unequal. The other three properties extend previous definitions of cooperation in symmetric social dilemmas[1,2] to the case of asymmetric public good games. In any game that satisfies **(IF)**, players have an incentive to contribute as little as possible. Without repeated interactions, such games have a unique Nash equilibrium according to which no player contributes, $\mathbf{x} = \mathbf{0}$.

We note that the four properties introduced above rule out certain threshold public good games, where payoffs suddenly increase when payoffs exceed a predefined threshold[3–8]. In such games, cooperation can already be sustained in one-shot interactions, because people have an incentive not to fall below the threshold. In contrast, here we are interested in games where the social dilemma is as stringent as possible, and where repeated interactions are key to maintain mutual cooperation.

In **Fig. 1** of the main text we have classified different public good games according to two dimensions, linearity and symmetry. A public good game is *linear* if a player's payoff is a linear function of her own endowment and of the absolute contributions of all other group members,

$$u_i(\mathbf{e}, \mathbf{x}) = \sum_{j}^{n} c_{ij} x_j e_j + e_i. \tag{1}$$

The constants c_{ij} determine to which extent player j's absolute contribution to the public good affects player i's payoff. We note that property **(PE)** requires $c_{ij} > 0$ for all co-players $j \neq i$, whereas prop-

3

erty **(IF)** requires $c_{ii} < 0$ for all players i. The property **(OC)** requires $c_{1j} + \ldots + c_{nj} > 0$ for all players j. Property **(C)** is satisfied since linear functions are continuous.

A public good game is *symmetric* if players do not differ in their personal characteristics, such that any payoff differences between players can entirely be explained by differences in their endowments and contributions. To formalize this notion of symmetry, let σ be a permutation of the numbers $(1, \ldots, n)$. For a given endowment vector $\mathbf{e}$ and contribution vector $\mathbf{x}$, let $\mathbf{e}_\sigma := (e_{\sigma(1)}, \ldots, e_{\sigma(n)})$ and $\mathbf{x}_\sigma := (x_{\sigma(1)}, \ldots, x_{\sigma(n)})$ be the permuted endowment and contribution vectors. We call the payoff function u symmetric if

$$u_i(\mathbf{e}, \mathbf{x}) = u_{\sigma(i)}(\mathbf{e}_\sigma, \mathbf{x}_\sigma) \quad \text{for all } \mathbf{e}, \mathbf{x}, \sigma. \tag{2}$$

That is, if players were to switch roles with respect to their endowments and their contributions, then their payoffs are switched accordingly. In particular, for symmetric payoff functions the group payoff is independent of the players' identities, $U(\mathbf{e}, \mathbf{x}) = U(\mathbf{e}_\sigma, \mathbf{x}_\sigma)$.

1.2 Repeated games

Herein, we consider individuals that do not interact once with each other, but repeatedly over several rounds. Specifically, we assume that after each round of the public goods game, there is another round with probability δ, with $0 < \delta < 1$. It follows that the expected number of rounds follows a geometric distribution with mean $1/(1-\delta)$. Player i receives the same endowment e_i in every round. However, the player's contribution x_i may change from one round to the next, and the player's decision how much to contribute in a given round may depend on previous outcomes of the game. Players make these decisions according to their strategies. A strategy for player i is a rule that tells the player how to choose her contribution $x_i(t)$ in round t, depending on the players' endowments $\mathbf{e}$ and depending on all players' previous contribution decisions.

For example, the strategy ALLD sets $x_{\text{ALLD}}(t) = 0$ in every round, independent of the players' endowment and independent of the outcome of previous rounds. Another example is the strategy Grim. A Grim player contributes her full endowment $x_{\text{Grim}}(0) = 1$ in the initial round. In all subsequent rounds, she continues to contribute her full endowment $x_{\text{Grim}}(t) = 1$, provided that all players have done so in all previous rounds. Otherwise, if at least one player has contributed less than his full endowment in some previous round, a Grim player contributes nothing, $x_{\text{Grim}}(t) = 0$.

In this way, the players' strategies determine which fraction of the endowment the players contribute in each round. If the contributions at time t are given by the vector $\mathbf{x}(t) = (x_1(t), \ldots, x_n(t))$, the players' payoffs for the repeated game are defined as the weighted average

$$\pi := (1 - \delta) \sum_{t=0}^{\infty} \delta^t u(\mathbf{e}, \mathbf{x}(t)). \tag{3}$$

We will sometimes consider the limiting case that there is always another round, and $\delta \to 1$. In that case,

the players' payoffs according to Eq. (3) approach the average payoff per round,

$$\pi = \lim_{T \to \infty} \frac{1}{T} \sum_{t=0}^{T-1} u(\mathbf{e}, \mathbf{x}(t)), \tag{4}$$

provided this limit exists. Herein, we will often consider players with bounded memory such that the existence of the limit in Eq. (4) is guaranteed[9].

2 Related literature

2.1 Related work in evolutionary game theory

Traditionally, most models on the evolution of cooperation study the behavior of individuals who have the same personal attributes, the same incentives to cooperate, and the same strategy set to choose from[10]. Only a few models have considered the effect of different sources of asymmetry. Here we provide a review of this previous work.

Bergstrom et al[11] consider a one-shot model of public good provision in which players decide how much of their wealth they want to spend either on a private or a public good. In their model, individual preferences are such that total contributions to the public good in equilibrium are always positive; in particular, mutual defection is not an equilibrium. Provided all players contribute, a sufficiently weak redistribution of wealth has no effect at all because players adjust their public good contributions accordingly. However, equalizing the initial wealth of players may decrease total contributions if it involves transfers from contributors to non-contributors.

The effect of asymmetric cooperation costs has also been studied in the context of the volunteer's dilemma[12,13]. In the volunteer's dilemma, it takes one group member to pay a cost for everyone in the group to derive some benefit. Asymmetric costs can facilitate coordination when players can make their choice in continuous time[13]: in equilibrium, the player with the lowest cost volunteers without delay.

Closer to our work, inequality in the players' endowments and their wealth have been considered in models of the collective risk-dilemma[8,14,15]. In the collective-risk dilemma, a group of players needs to collect a fixed target sum in order to avoid dangerous climate change[3]. If they fail to meet the required threshold, they lose their remaining endowment with some exogenous loss probability. Vasconcelos et al[14] consider a model in which players are either 'rich' or 'poor', and in which players can decide whether or not they wish to contribute a fixed fraction of their endowment to the public pool. In this model, heterogeneity can promote cooperation, but only in cases when poor players imitate the strategies of the rich (and vice versa). In Abou Chakra et al[8], players differ along two dimensions, their wealth and their loss probability. Players who have more to lose – either because they are wealthier or because they have a higher loss probability – tend to make higher contributions.

Finally, Akin[16] considers memory-1 strategies in the asymmetric repeated prisoner's dilemma without discounting. Building on the recently developed theory of zero-determinant strategies[17–22], he characterizes all strategies that can be used to sustain mutual cooperation.

5

We add to this evolutionary game theory literature in the following way:

(*i*) We provide a general framework that is applicable to social dilemmas with an arbitrary number of players and continuous degrees of cooperation.

(*ii*) We allow for various different dimensions of heterogeneity. For example, players can simultaneously differ in their endowments, productivities, or their benefit from the public good. We predict how these different forms of inequality interact.

(*iii*) We explore direct reciprocity in strict social dilemmas, while most of the previous work had its focus on inequality in one-shot games, or in games with a finite horizon.

(*iv*) In our applications, the threat of future retaliation is key to sustain cooperation. The logic of reciprocity seems to suggest that cooperation is most likely if all players have the same endowment, yet we show that this does not need to be the case.

(*v*) We employ three different, but complementary approaches: equilibrium calculations, individual-based simulations, and a behavioral experiment. Each approach is based on slightly different assumptions, yet they all yield the same qualitative results.

2.2 Related experimental work

The public goods game is one of the classical paradigms to study the dynamics of cooperation in behavioral experiments. The vast literature on this topic suggests that individual contributions depend on many different aspects, including the net cost of contributions, communication opportunities, and thresholds, see the reviews of Ledyard[23] and Zelmer[24]. In the following, we summarize previous experimental work in which participants either experienced only one kind of inequality, or in which different kinds of inequality interacted.

Effect of one-dimensional inequality. Several studies have explored the origins and consequences of endowment inequality in public good games. For example, Buckley and Croson[25] study a linear public goods game in groups in which half of the players have 25 tokens per round, whereas the other half has 50 tokens per round. They find that both kinds of players contribute roughly the same absolute amount (that is, players with a higher endowment make a lower relative contribution). In a similar experiment, Hargreaves Heap et al[26] compare the contributions of homogeneous and heterogeneous groups. Their participants either have 20, 50, or 80 tokens, and they interact in groups of size 3. Players who are either given 20 or 50 tokens play consistently, independent of whether they are in a homogeneous or in a heterogeneous group. However, 80-token players contribute significantly less in mixed groups, compared to homogeneous groups in which everyone receives 80 tokens. In the experiment of Cherry et al[27], players either earn their initial endowments in a quiz, or the endowments are randomly assigned by the experimenter. This study finds that the origin of the endowment has no effect, whereas heterogeneity decreases the players' contributions. Similarly, in the experiment of Kesternich et al[28], subjects are either assigned random endowments, or they earn their endowments by outperforming their peers in a 'slider

task' that requires effort. Again, when contributions are voluntary, the origin of endowment seems to have a minor effect, although highly endowed subjects tend to contribute less if they have earned their endowment in the effort task.

Gächter et al[29] explore the consequences of endogenous inequality. In their linear public goods game, all players are equipped with the same starting capital. However, a player's accumulated wealth in round t is his endowment for round $t+1$. The study finds that inequality (as measured by the Gini coefficient) is increasing over time. Moreover, there is substantial variation in inequality between groups. In a similar experiment, Cardigan et al[30] study contributions in public good games when a player's payoff in odd rounds serves as the player's endowment in the subsequent even rounds. They find that under certain conditions, this "carryover" can increase contributions in odd rounds. Nishi et al[31] consider pairwise prisoner's dilemma games among players on a dynamical network. Their treatments differ in the degree of inequality, and in whether players are informed about their co-player's accumulated wealth. As their main result they find that heterogeneity has a negligible impact on cooperation per se, but visible wealth inequality has a negative impact[32].

Moreover, several studies have explored how heterogeneity interacts with different mechanisms for the maintenance of cooperation. These mechanisms include communication[33,34], voting[35,36], institutional sharing rules[28,37], mandatory minimum contributions[38], or sanctions[29,36,39]. For example, the experiment by Reuben and Riedl[39] explores the interaction of endowment inequality and costly punishment. There are $2 \times 4 = 8$ treatments. The treatments differ in whether or not subjects are able to punish each other after each round, and in the type of endowment or benefit inequality considered. This study finds that without punishment, contributions steadily decline across all groups, such that any difference between groups is blurred. In contrast, when punishment is possible, contributions are high in all group types. Moreover, depending on the type of inequality, different contribution norms are enforced.

Without any additional mechanisms, the existing data suggests that unequal endowments tend to decrease contributions, whereas the effect of unequal productivities is ambiguous. Our own experiment agrees with this view: when the two players are equally productive, unequal endowments diminish overall cooperation. Conversely, when players have equal endowments, differences in the players' productivity have a negligible effect on contributions. Interestingly, in most of the aforementioned experiments on asymmetric public good games, the number of rounds is fixed and known to the participants. For such repeated social dilemmas with a predefined horizon, backward induction suggests that participants should not cooperate at all. In contrast to this prediction, a complete unravelling of contributions is rarely observed in repeated social dilemma experiments[40]. In such cases, our theoretical framework for indefinitely repeated games may provide a more accurate prediction of the qualitative behavior in asymmetric public good games, even if subjects are informed about the exact length of the game ex ante.

Effect of interacting inequalities. The interaction of different forms of inequality has received comparably little attention, with two exceptions. First, in one treatment of Chan et al[33], the authors allow for two sources of heterogeneity. In that treatment, players could differ in their endowments and in the

7

benefit they derived from the public good (the treatment only allows for what we would call aligned inequality). The authors use a non-linear payoff function, and parameters have been chosen such that mutual defection cannot arise as an equilibrium. As a consequence, their experiment does not meet the conditions of a strict social dilemma, as discussed in Section 1. Nevertheless, their results are in line with our prediction: they find that in the absence of communication, aligned inequality can help cooperation.

Second, van Gerwen et al[41] explore cooperative behavior in a contextualized public goods game between three subjects. Subjects may differ in their 'skills' (akin to our endowment), and in whether or not they have received 'training' (akin to our differences in productivities). In addition, groups may differ in how many rounds subjects interact with each other. Groups either remain stable for six consecutive rounds ('stable teams'), or they are formed anew after each round ('altering teams'). For comparison with our study, which is based on direct reciprocity, the stable team treatments of van Gerwen et al are most relevant. For those treatments, the authors find that more skilled subjects (i.e., players with higher endowment) tend to make higher absolute contributions. Interestingly, however, a player's contribution is neither affected by his training, nor is there a significant interaction between training and skills. Moreover, on the group level van Gerwen et al find no difference in contributions based on whether training is provided to high-skill or low-skill subjects (see the 'High' and 'Low' condition in their Figure 2). That is, they find no difference between aligned and misaligned inequality as we do. There are two reasons that might explain why their conclusions differ:

1. *The comparably short duration of their experiment:* Their stable teams only interact for six rounds. Moreover, the total number of rounds is commonly known, leading to a significant end-game effect (as confirmed by their Table 4). The horizon of the stable team experiments may thus be too short for reciprocity to unfold.

2. *The framing of their experiment:* To increase external validity, subjects in van Gerwen et al are provided with strongly contextualized instructions. For example, subjects learn that their training is a gift by the employer. The authors themselves expect that these contextualized instructions increase the contributions of those individuals who receive training. Such framing effects might explain why groups with aligned inequality did not outperform groups with misaligned inequality. The framing might also explain why some statistical relationships became significant when altering teams were considered instead of stable teams.

In sum, we add to the previous experimental literature on public good games in the following way.

(*i*) We provide a general theoretical framework to explain reciprocal behavior in repeated social dilemmas. Our framework applies to arbitrary public good games. Players may be symmetric or asymmetric, and their payoffs may be linear or nonlinear.

(*ii*) We experimentally confirm previous results on the isolated effects of endowment and productivity inequality.

(*iii*) We are first to show that in repeated games, endowment and productivity inequality interact nontrivially. Aligned inequality can enhance cooperation, whereas misaligned inequality reduces it.

3 Equilibrium analysis

Whereas one-shot public good games only allow for equilibria with zero contributions, repeated games can give rise to alternative equilibria[9,42]. In the following, we are particularly interested in the subgame perfect equilibria in which players contribute their full endowment in every round. A strategy profile is called a subgame perfect equilibrium if, after any previous history of play, no player has an incentive to deviate from her strategy[43]. The subgame perfect equilibrium concept represents an equilibrium refinement of the Nash equilibrium: every subgame perfect equilibrium is also a Nash equilibrium, while the converse does not need to hold[44].

3.1 Feasibility of cooperation

We say that full cooperation is feasible in a given public good game if there is a subgame perfect equilibrium in which all players contribute their full endowment in every round. For a given payoff function u and continuation probability δ, let $E_u(\delta)$ denote the set of all endowment distributions that allow for full cooperation,

$$E_u(\delta) = \Big\{ \mathbf{e} = (e_1, \dots, e_n) \,\Big|\, \text{Full cooperation is feasible for the given endowment distribution } \mathbf{e} \Big\}. \tag{5}$$

In the upper triangles in **Fig. 2**, the respective sets $E_u(\delta)$ are represented by blue-shaded areas. In the following Proposition, we give a simple characterization of the set $E_u(\delta)$ for arbitrary public good games. We use the notation $\mathbf{x} = \mathbf{1}_{-i}$ to refer to the contribution vector where all players but player i give their full endowment, whereas player i contributes nothing.

Proposition 1. *Consider a public good game with payoff function u and continuation probability $\delta < 1$, and let $\mathbf{e}$ denote some endowment distribution. The following are equivalent:*

1. *$\mathbf{e} \in E_u(\delta)$*
2. *The following condition holds for all players i with $e_i > 0$,*

$$\delta \Big(u_i(\mathbf{e}, \mathbf{1}_{-i}) - u_i(\mathbf{e}, \mathbf{0}) \Big) \geq u_i(\mathbf{e}, \mathbf{1}_{-i}) - u_i(\mathbf{e}, \mathbf{1}) \tag{6}$$

3. *The strategy profile where all players apply the strategy Grim is a subgame perfect equilibrium for the given endowment distribution $\mathbf{e}$.*

All proofs for the results in this section are in the **Appendix**. The above Proposition says that if full cooperation can be sustained at all for a given endowment distribution $\mathbf{e}$, then it can be sustained with the strategy Grim. In most applications, the stability of Grim can be checked easily. Using Proposition 1, we can show that $E_u(\delta)$ exhibits the usual monotonic dependence on the continuation probability δ.

Proposition 2. *Consider a public good game with payoff function u.*

1. *Suppose δ and δ' are two continuation probabilities with $\delta < \delta'$. Then $E_u(\delta) \subset E_u(\delta')$.*

2. *There is a $\delta' < 1$ such that $E_u(\delta) \neq \emptyset$ for all $\delta \geq \delta'$.*

That is, the set of all endowment distributions for which cooperation is feasible increases with the continuation probability. If a given endowment distribution allows for full cooperation for a given continuation probability, then it also allows for full cooperation for all larger continuation probabilities. Moreover, cooperation is always feasible if only the continuation probability δ is sufficiently large.

Using Proposition 1 we can also show that cooperation is never feasible if initial endowments are too strongly biased in favor of one player.

Proposition 3. *Consider an arbitrary public good game with continuation probability $\delta < 1$ and endowment distribution $\mathbf{e}$. For any player i, there is an $\varepsilon_i > 0$ such that if player i's endowment exceeds $e_i = 1 - \varepsilon_i$ then full cooperation is infeasible.*

The result of Proposition 3 also becomes apparent in **Fig. 2** of the main text. In all triangles depicted, there is a neighborhood around the corners of the triangle where cooperation is not feasible. Intuitively, full cooperation in repeated games can only be sustained if any deviation in the present round can be punished by withholding contributions in future rounds. However, if the initial endowments are too concentrated such that one player is given an excessive share, her co-players cease to have enough leverage to prevent her from free riding.

It may appear that a similar argument could also show that an equal endowment distribution is generally most conducive to cooperation. In the following sections we show that this conjecture is true when the public good game is linear and symmetric, whereas it is false if payoffs are nonlinear or asymmetric.

3.2 Linear and symmetric public good games

A linear public goods game (1) satisfies the symmetry condition (2) if and only if the payoff function has the form

$$u_i(\mathbf{e}, \mathbf{x}) = c_1 \sum_{j \neq i} x_j e_j - c_2 x_i e_i + e_i. \tag{7}$$

Here, the constant c_1 measures by how much player i's payoff is affected by co-player j's contribution (due to symmetry, this constant is independent of both i and j). Similarly, c_2 measures to which extent player i's payoff depends on her own contribution; again, due to symmetry c_2 is independent of i. The three properties **(PE)**, **(IF)**, and **(OC)** then imply $c_1 > 0$, $c_2 > 0$ and $(n-1)c_1 - c_2 > 0$, respectively. In the main text we have considered scenarios where total contributions of all players are multiplied by a joint productivity factor r, with $1 < r < n$, and then equally shared among all group members. This standard formulation of public goods games can be recovered as a special case of game (7) when we set $c_1 = r/n$ and $c_2 = 1 - r/n$.

By applying Proposition 1 to the linear and symmetric public goods game (7), we find that cooperation is feasible for a given endowment distribution $\mathbf{e}$ if and only if

$$\delta c_1 (1 - e_i) \geq c_2 e_i \quad \text{for all players } i \text{ with } e_i > 0. \tag{8}$$

10

Equivalently, cooperation is feasible for the endowment distribution $\mathbf{e}$ if and only if

$$\delta \frac{c_1}{c_2} \geq \frac{\max\{e_i\}}{1 - \max\{e_i\}}. \tag{9}$$

Condition (9) indicates that it is always the player with the highest initial endowment who has the largest temptation to deviate from full cooperation.

Now, suppose that cooperation is feasible for some endowment distribution, such that $E_u(\delta) \neq \emptyset$. Thus, there needs to be an endowment distribution $\mathbf{e} = (e_1, \ldots, e_n)$ that satisfies inequality (9). But since $\max\{e_i\} \geq 1/n$ and since the map $z \mapsto z/(1-z)$ is monotonically increasing for $0 < z < 1$, it then follows that also the equal endowment distribution $(1/n, \ldots, 1/n)$ satisfies inequality (9). That is, if full cooperation is feasible at all in some given linear and symmetric public good game, then it can always be achieved by assigning equal endowments to all players.

By plugging the equal endowment distribution $\mathbf{e} = (1/n, \ldots, 1/n)$ into condition (9), we can conclude that cooperation is feasible in a linear and symmetric public goods game if and only if $\delta \geq \delta^*$ with $\delta^* := c_2/((n-1)c_1)$. In particular, since $(n-1)c_1 - c_2 > 0$ it follows that $\delta^* < 1$, and thus cooperation is always feasible if the public goods game is repeated sufficiently often. We can summarize these observations as follows.

Proposition 4. *Consider a linear and symmetric public goods game with payoff function (7).*

1. *If full cooperation is feasible for some arbitrary endowment distribution, it is also feasible when players have equal endowments. That is, if $E_u(\delta) \neq \emptyset$ for some δ, then $(1/n, \ldots, 1/n) \in E_u(\delta)$.*
2. *Full cooperation is feasible, that is $E_u(\delta) \neq \emptyset$, if and only if $\delta \geq \dfrac{c_2}{c_1} \dfrac{1}{n-1}$.*

This first result in Proposition 4 is also apparent in **Fig. 2a**: the center of the triangle which represents the point of equal endowments is always an element of the blue area where full cooperation is feasible, provided δ is sufficiently large for this blue area to exist.

3.3 Asymmetric linear public good games

As an example of a public goods game that is linear but asymmetric, we have considered the following example in the main text,

$$u_i(\mathbf{e}, \mathbf{x}) = \frac{1}{n} \sum_j r_j x_j e_j + (1 - x_i) e_i. \tag{10}$$

We have interpreted r_j as the productivity of player j's contributions to the public good. The payoff function (10) satisfies the two properties **(IF)** and **(OC)** if and only if the individual productivities satisfy $1 < r_j < n$. If that is the case, property **(PE)** holds automatically, whereas property **(C)** is trivially satisfied for all public good games of the form (10).

When we apply Proposition 1 to the asymmetric public goods game (10), we yield the following

necessary and sufficient condition for the feasibility of cooperation:

$$\frac{\delta}{n} \sum_{j \neq i} r_j e_j \geq \left(1 - \frac{r_i}{n}\right) e_i \quad \text{for all players } i \text{ with } e_i > 0. \tag{11}$$

In particular, for equal endowments, $\mathbf{e} = (1/n, \ldots, 1/n)$, full cooperation is feasible if and only if

$$\delta \geq \frac{n - \min(r_j)}{\sum_{j=1}^{n} r_j - \min(r_j)}. \tag{12}$$

Condition (12) indicates that when players are given equal endowments, the player with the lowest productivity will face the largest temptation to deviate (because this player faces the highest marginal cost $1 - r_i/n$ of contributing). This observation suggests that when players differ in their productivities, cooperation may be easier to sustain when players also differ in their endowments.

In **Fig. 2b** of the main text, we have provided a simple example. There, we have considered a 3-player public goods game with $r_1 = 2.9$, $r_2 = 1.5$, $r_3 = 1.1$, and continuation probability $\delta = 0.3$. A straightforward computation of the conditions (11) shows that when players have equal endowment, full cooperation is infeasible, because both players 2 and 3 (the players with comparably low productivity) have an incentive to deviate. In contrast, if we redistribute endowments such that $\mathbf{e} = (4/5, 1/5, 0)$, the feasibility condition (11) is satisfied for all players. Interestingly, according to this distribution most of the initial endowment goes to the most productive player 1, and player 3 does not receive any endowment at all. However, player 2's share of the endowment, $e_2 = 1/5$, is sufficient to still incentivize player 1 to cooperate.

Moreover, we note that even if full cooperation were feasible with equal endowments, players may still fare better when a higher share of the endowment is given to the more productive player. To illustrate this point, consider the above public good game with $r_1 = 2.9$, $r_2 = 1.5$, $r_3 = 1.1$, but with the continuation probability increased to $\delta = 0.5$. In that case, cooperation is also feasible for an equal endowment distribution. However, in that equilibrium, each player's payoff is $(r_1 + r_2 + r_3)/9 \approx 0.61$, which is smaller than the mutual cooperation payoff for $\mathbf{e} = (4/5, 1/5, 0)$, yielding a payoff of $(4r_1 + r_2)/15 \approx 0.87$. We can summarize these observations as follows.

Proposition 5.

1. *There are asymmetric linear public good games such that $E_u(\delta) \neq \emptyset$, but $(1/n, \ldots, 1/n) \notin E_u(\delta)$.*
2. *There are asymmetric linear public good games such that even if full cooperation is feasible for equal endowments, unequal endowments allow for higher equilibrium payoffs for every player.*

While the above Proposition shows that some inequality can improve the stability of cooperation, it does not clarify which kind of inequality is most conducive to cooperation. The following result shows that the conditions for the feasibility of cooperation are most easily satisfied if the players' endowments match their productivities. The proof is given in the **Appendix**.

12

Proposition 6. *Consider the public goods game with payoff function (10), and assume without loss of generality that players are ordered according to their productivities, $r_1 \leq r_2 \leq \ldots \leq r_n$. Let $\mathbf{e}' = (e'_1, \ldots, e'_n)$ be some endowment vector, and let $\mathbf{e}'' = (e''_1, \ldots, e''_n)$ be the permuted endowment vector such that the ordering of the entries in $\mathbf{e}''$ is aligned with the players' productivities; that is, $\mathbf{e}'' = \mathbf{e}'_\sigma$ for some permutation σ and $e''_1 \leq e''_2 \leq \ldots \leq e''_n$. If $\mathbf{e}' \in E_u(\delta)$, then $\mathbf{e}'' \in E_u(\delta)$.*

Proposition 6 states that if full cooperation can be sustained for some endowment distribution, then it can always be sustained by redistributing the players' endowments such that more productive players have higher endowments.

So far, we have been concerned with asymmetric games in which players have differed in their productivity. Alternatively, we may also consider games in which players are equally productive, but they differ in the extent to which they benefit from publicly provided goods. Such games may be represented by the payoff function

$$u_i(\mathbf{e}, \mathbf{x}) = q_i \cdot r \sum_j x_j e_j + (1 - x_i) e_i. \tag{13}$$

The coefficient r with $1 < r < n$ represents the players' common productivity. The coefficient q_i determines the relative value of the public good to player i. Without loss of generality, we may assume that these coefficients are normalized such that $q_1 + \ldots + q_n = 1$. In order for the game to qualify as a social dilemma, **(PE)** requires that $q_i > 0$ for all players i, whereas **(IF)** requires $q_i < 1/r$.

Again, it is straightforward to construct examples of a public good game with payoff function (13) such that full cooperation can only be sustained for unequal endowment distributions. According to Proposition 1, cooperation is feasible for a given endowment distribution $\mathbf{e}$ if and only if

$$\delta q_i r (1 - e_i) \geq (1 - q_i r) e_i \quad \text{for all players } i \text{ with } e_i > 0. \tag{14}$$

For equal endowments, $\mathbf{e} = (1/n, \ldots, 1/n)$, it follows that full cooperation is feasible if and only if

$$\delta \geq \frac{1 - r \cdot \min(q_i)}{(n-1) r \cdot \min(q_i)}. \tag{15}$$

Similar to before, under equal endowments it is the player with the highest marginal cost to contribute who is most likely to deviate (here, the player with the smallest benefit from the public good). To prevent this player from deviating, it may again be effective to redistribute initial endowments, such that players with a stronger incentive to contribute get a larger share of the endowment.

As a numerical example, let us consider a 3-player public goods game with $q_1 = 0.1$, $q_2 = 0.3$, $q_3 = 0.6$, $r = 1.5$ and $\delta = 0.5$. A straightforward computation shows that when players have equal endowments, condition (14) is violated for players 1 and 2. Instead, the inequalities in (14) hold for all players if endowments are redistributed such that $\mathbf{e} = (0, 1/5, 4/5)$.

We note, however, that games in which players derive unequal benefits from the public good differ in

two important aspects from games where players have unequal productivities. First, when players derive unequal benefits, player i's payoff in the fully cooperative equilibrium is $\pi_i = q_i r$, and thus independent of the initial endowments. That is, provided players still achieve full cooperation, a redistribution of endowments has no payoff effects. Second, in the 3-player game with different productivities, everyone preferred the mutual cooperation payoff for the endowment distribution $(0, 1/5, 4/5)$ to the mutual defection payoff when players have equal endowments. In the 3-player game where players derive different benefits from the public good, this does not need to be the case anymore. In the example discussed in the previous paragraph, player 1 yields a payoff of $1/3$ if all players are given an equal endowment and everyone defects. In contrast, player 1 yields the lower payoff of 0.15 if the initial endowment distribution is $\mathbf{e} = (0, 1/5, 4/5)$ and everyone cooperates. That is, even if a redistribution of endowments from $(1/3, 1/3, 1/3)$ to $(0, 1/5, 4/5)$ favors mutual cooperation, players may no longer unanimously agree to such a redistribution.

3.4 An example of a nonlinear symmetric public goods game

As a simple example of a nonlinear but symmetric public goods game, **Fig. 2c,f** of the main text illustrates the case of a game with payoff function

$$u_i(\mathbf{e}, \mathbf{x}) = c_1 \cdot \max_{j,\, k \neq j} \left(e_j x_j + e_k x_k \right) + c_2 \cdot \sum_{j=1}^{n} e_j x_j + (1 - x_i) e_i, \tag{16}$$

for constants $c_1, c_2 > 0$. Due to the first term on the right hand's side, the payoff function u is piecewise linear but not linear for $n \geq 3$. This game reflects a situation in which the two highest absolute contributions are particularly important for the success of the public good. Payoff function (16) satisfies property **(IF)** if $c_1 + c_2 < 1$, and it satisfies property **(OC)** if $c_2 > 1/n$.

Given that the payoff function (16) provides an extra benefit to the two highest absolute contributions, an equal distribution of endowments may no longer be optimal for groups of size $n \geq 3$. According to Proposition 1, full cooperation can be sustained for equal endowments if and only if

$$\delta \geq \frac{1 - c_2}{2c_1 + (n-1)c_2}. \tag{17}$$

When players have unequal endowments, and without loss of generality $e_1 \geq e_2 \geq \ldots \geq e_n$, Proposition 1 yields the following condition for full cooperation,

$$\delta \geq \max \left\{ \frac{(1 - c_1 - c_2)e_1 + c_1 e_3}{c_1(e_2 + e_3) + c_2(1 - e_1)}, \frac{(1 - c_2)e_3}{c_1(e_1 + e_2) + c_2(1 - e_3)} \right\}. \tag{18}$$

Condition (18) indicates that if any player has an incentive to deviate from full cooperation, it is either the player with the highest endowment, or the player with the third-highest endowment (i.e., the first player whose contribution is not multiplied by the factor c_1). By comparing the conditions (17) and (18), we obtain the following result.

14

Proposition 7. *Consider a nonlinear public good game between $n \geq 3$ players with payoff function (16). Let $\mathbf{e}'$ be the equal endowment distribution. Let $\mathbf{e}''$ be a distribution where the full endowment is equally split among two players only; that is, there are players i, j such that $e_i'' = e_j'' = 1/2$ and $e_k'' = 0$ for all other k. Then for all continuation probabilities $\frac{1-c_1-c_2}{c_1+c_2} \leq \delta < \frac{1-c_2}{2c_1+(n-1)c_2}$ we have $\mathbf{e}' \notin E_u(\delta)$, $\mathbf{e}'' \in E_u(\delta)$.*

The content of Proposition 7 is illustrated in **Fig. 2c**, where we have considered the case $n = 3$ with $c_1 = 1/2$, $c_2 = 1/3$. For these parameter values, equal endowments can maintain cooperation for $\delta \geq 0.4$, whereas splitting the endowment between two players can maintain cooperation for $\delta \geq 0.2$. For the figure, we have used the value $\delta = 0.35$; in that case, cooperation is feasible in three separate regions, centered around the three points where two players equally split the full endowment.

3.5 Maximally cooperative endowment distributions

In the remainder of this section, we wish to introduce a notion for endowment distributions that are most conducive to cooperation. To this end, let us first formally define the minimal continuation probability compatible with full cooperation. For a given public good game with payoff function u, we define this minimal continuation probability δ_u^* as

$$\delta_u^* := \inf \left\{ \delta \in [0, 1] \mid E_u(\delta) \neq \emptyset \right\}. \tag{19}$$

Because of Proposition 2, this minimal continuation probability is well-defined and satisfies $0 < \delta_u^* < 1$. Using this minimal continuation probability we introduce the following notion of maximally cooperative endowment distributions.

Definition. *For a given public good game u we say that an endowment distribution $\mathbf{e}^* = (e_1^*, \dots, e_n^*)$ is maximally cooperative if $\mathbf{e}^* \in E_u(\delta_u^*)$.*

That is, an endowment distribution $\mathbf{e}^*$ is maximally cooperative if it requires the lowest continuation probability for cooperation to be feasible. Due to Proposition 2, it follows that such a distribution is most conducive to cooperation in the following sense: If cooperation is feasible for *any* endowment distribution for a given δ, then it is guaranteed to be feasible when players receive the maximally cooperative endowment distribution $\mathbf{e}^*$.

To find the maximally cooperative endowment distribution, one needs to identify a pair $(\mathbf{e}^*, \delta^*)$ such that all inequalities (6) are satisfied; moreover, for any other such pair $(\mathbf{e}, \delta)$ it needs to be true that $\delta \geq \delta^*$. Because the left hand's side of the inequalities (6) may involve products of e_i and δ, this generally yields a nonlinear optimization problem. In the following we solve this optimization problem for the special case of a pairwise game with a linear payoff function.

Proposition 8. *Consider a public good game between two players with linear payoff function u given*

15

by Eq. (1)*, such that*

$$u_1(\mathbf{e}, \mathbf{x}) = c_{11}x_1e_1 + c_{12}x_2e_2 + e_1$$
$$u_2(\mathbf{e}, \mathbf{x}) = c_{21}x_1e_1 + c_{22}x_2e_2 + e_2, \tag{20}$$

for constants $c_{11}, c_{22} < 0$ and $c_{12}, c_{21} > 0$. Then the minimal continuation probability δ_u^* and the maximally cooperative endowment distribution $\mathbf{e}^* = (e_1^*, e_2^*)$ are given by

$$\delta_u^* = \sqrt{\frac{c_{11}c_{22}}{c_{12}c_{21}}} \quad \text{and} \quad \frac{e_1^*}{e_2^*} = \sqrt{\frac{c_{12}c_{22}}{c_{11}c_{21}}}. \tag{21}$$

In the main text, we have considered the example where players have different productivities, such that $c_{ij} = r_j/n$ for $j \neq i$ and $c_{ii} = r_i/n - 1$. In that special case, the maximally cooperative endowment distribution according to Eq. (21) becomes

$$\frac{e_1}{e_2} = \sqrt{\frac{r_2(2 - r_2)}{r_1(2 - r_1)}}. \tag{22}$$

In particular, it follows that an equal endowment distribution $e_1 = e_2$ is maximally cooperative if and only if the two players coincide in their productivities. If players differ in their productivities, the more productive player is supposed to yield a higher endowment according to the maximally cooperative endowment distribution.

Similarly, for the game with asymmetric benefits considered in Section 3.3, we can set $c_{ij} = q_i r/n$ and $c_{ii} = q_i r/n - 1$. In this case, the maximally cooperative endowment distribution becomes

$$\frac{e_1}{e_2} = \sqrt{\frac{q_1(2 - rq_2)}{q_2(2 - rq_1)}}. \tag{23}$$

Again, we conclude that equal endowments are only maximally cooperative if $q_1 = q_2$, such that both players derive the same benefit from the public good. Otherwise, if there is a player who derives a larger benefit, this player is supposed to receive the larger endowment.

4 Evolutionary analysis

Whereas the previous results have clarified when full cooperation is feasible, based on static equilibrium considerations, here we explore whether cooperation can actually evolve if players have not yet settled on a particular equilibrium. To make such an evolutionary approach computationally tractable, we first focus on a a reduced strategy space.

4.1 Memory-one strategies

Compared to the previous equilibrium analysis, the strategy space considered in the following is based on two simplifying assumptions. First, we assume that in each round, players only choose among a finite

set X of contribution levels, $X = \{\hat{x}_1, \ldots, \hat{x}_m\}$. In the most simple case that $X = \{0, 1\}$, we refer to the two possible choices as 'defection' (D) and 'cooperation' (C), respectively. Second, we assume that players use *memory-one* strategies: a player's decision to cooperate in the next round may depend on the outcome of the previous round, but it is independent of all earlier rounds. Memory-one strategies have the major advantage that the payoffs according to Eqs. (3) and (4) can be computed efficiently, as we describe further below.

Memory-one strategies can be represented by a vector $\mathbf{p} = (p_{0,k}, \; p_{\mathbf{x},k})$, with entries in the unit interval $[0, 1]$. The entries $p_{0,k}$ for $1 \leq k \leq m$ denote the probability that in the first round, the player contributes a fraction $\hat{x}_k$ of the endowment . The remaining entries $p_{\mathbf{x},k}$ encode the conditional probability that the player contributes $\hat{x}_k$ in any subsequent round, given that the contribution vector of the previous round was $\mathbf{x} \in X^n$. For the vector $\mathbf{p}$ to be a sensible strategy, we require that in each round, the respective probabilities sum up to one, $\sum_k p_{0,k} = 1$ and $\sum_k p_{\mathbf{x},k} = 1$ for all contribution vectors $\mathbf{x} \in X^n$. If all entries of $\mathbf{p}$ are either equal to 0 or equal to 1, we say the memory-one strategy is *deterministic* or *pure*. Otherwise the strategy is called *stochastic*.

When all n players use memory-one strategies $\mathbf{p}^1, \ldots, \mathbf{p}^n$, the resulting payoffs can be calculated using a Markov chain approach. The states of this Markov chain are the m^n possible contribution vectors $\mathbf{x} \in X^n$ that determine the outcome of a given round. If $\mathbf{x}$ is the contribution vector of the previous round, the probability that the next round's contribution vector is $\mathbf{x}' = \{\hat{x}_{k_1}, \ldots, \hat{x}_{k_n}\}$ can be calculated as $m_{\mathbf{x},\mathbf{x}'} = p_{\mathbf{x},k_1}^1 \cdot \ldots \cdot p_{\mathbf{x},k_n}^n$. The corresponding transition probabilities can be assembled in a matrix $M = (m_{\mathbf{x},\mathbf{x}'})$. In addition, we need to specify the probability to observe each contribution vector $\mathbf{x}' = \{\hat{x}_{k_1}, \ldots, \hat{x}_{k_n}\}$ in the very first round. This probability is given by $v_{\mathbf{x}'}^0 = p_{0,k_1}^1 \cdot \ldots \cdot p_{0,k_n}^n$. Given the initial distribution $\mathbf{v}^0 = (v_{\mathbf{x}}^0)$ and the transition matrix $M = (m_{\mathbf{x},\mathbf{x}'})$, we calculate

$$\mathbf{v} = (1 - \delta)\mathbf{v}^0 \cdot (I - \delta M)^{-1}, \tag{24}$$

where I is the $m^n \times m^n$ identity matrix. The entries of the vector $\mathbf{v} = (v_{\mathbf{x}})$ give the weighted probability to observe each possible outcome $\mathbf{x} \in X^n$ over the course of the repeated game. For $\delta \to 1$, the vector $\mathbf{v}$ approaches a left eigenvector of M with respect to the eigenvalue 1. Given the vector $\mathbf{v}$, we can calculate the players' payoffs according to Eqs. (3) and (4) as

$$\pi = \sum_{\mathbf{x} \in X^n} v_{\mathbf{x}} \cdot u(\mathbf{e}, \mathbf{x}). \tag{25}$$

For our evolutionary analysis, we suppose that the players' decisions can be subject to noise. That is, in any given round there is some small probability $\varepsilon > 0$ that a player implements a wrong move. For example, a player might contribute less than he actually wanted, or he may cooperate although he actually wanted to defect. As a consequence, if the player's actual memory-1 strategy is $\mathbf{p}$, his effective strategy becomes $(1 - \varepsilon) \cdot \mathbf{p} + \varepsilon/(m - 1) \cdot (\mathbf{1} - \mathbf{p})$. This assumption ensures that the transition matrix M is positive (that is, each entry of M is positive). In the limit $\delta \to 1$, the presence of noise thus implies that the invariant distribution $\mathbf{v}$ is independent of the outcome of the initial round. Qualitatively, the

evolutionary dynamics is largely independent of the exact value of $\varepsilon > 0$, provided errors are sufficiently rare (**Extended Data Fig. 4**).

4.2 A model of introspection dynamics

In the following we introduce a simple framework to model how players adopt new memory-one strategies over time. When games are symmetric, it is common to consider a process where players tend to imitate strategies of more successful players[45], or where successful players produce more offspring[46]. For games among unequals, however, such processes may yield counterintuitive results, because a strategy that is successful for one player may be disadvantageous for another. Alternatively, to study the dynamics in asymmetric matrix games, past work has considered players who are drawn from different populations, and who only adopt strategies that are successful within their own population[47,48]. However, for games in large groups such a multi-population approach seems unrealistic. In many of our applications, different players may interact with each other although they do not stem from different well-defined populations. To model the strategy dynamics among unequal players, we thus introduce a novel evolutionary process. Our process is based on the idea that occasionally, players may come up with new strategies, and that they compare their payoff so far with the payoff they could have got if they had used the new strategy all along. We call this process *introspection dynamics*.

Specifically, in a group of n players, we assume that in each evolutionary time step one of the players is chosen at random. This player i then picks a random new strategy of her available strategy set (while the co-players' strategies remain fixed). If the player's previous strategy yielded a payoff π_i, whereas the new strategy would have yielded the payoff $\tilde{\pi}_i$, we assume that the player switches to the new strategy with probability

$$\rho_s(\pi_i, \tilde{\pi}_i) = \frac{1}{1 + \exp\left[-s(\tilde{\pi}_i - \pi_i)\right]}. \tag{26}$$

The parameter $s \geq 0$ reflects the strength of selection. In the boundary case $s = 0$, the switching probability simplifies to $\rho = 1/2$, such that strategy changes occur at random. In the limit of strong selection, $s \to \infty$, the player only adopts the new strategy if it yields at least the payoff of the old strategy.

In **Figs. 2 – 4** of the main text, we have shown simulation results for this introspection dynamics. There, we have assumed that initially all players use ALLD. Then we have iterated the above described elementary updating process, assuming that subjects can choose among all available memory-1 strategies and that payoffs are computed by Eq. (25). By iterating this process over many time steps t, we yield a sequence of strategies $\left(\mathbf{p}^1(t), \ldots, \mathbf{p}^n(t)\right)$, where $\mathbf{p}^i(t)$ denotes the memory-1 strategy employed by player i after t time steps. Based on this sequence, we can calculate the average cooperation rates and payoffs of all players over the entire simulation. The respective MATLAB implementation is provided in the **Appendix**.

4.3 Introspection dynamics for players with finitely many strategies

If players employ pure memory-one strategies, such that the total number of strategies k is finite, we can use an exact method to compute expected trajectories of the introspection dynamics. The introspection dynamics takes the form of a discrete Markov chain with k^n possible states $P = (\mathbf{p}^1, \ldots, \mathbf{p}^n)$. The transition probability from state P to state $\tilde{P}$ is given by

$$
w_{P,\tilde{P}} = \begin{cases} \dfrac{1}{n(k-1)} \cdot \rho_s(\pi_i, \tilde{\pi}_i) & \text{if } \mathbf{p}^i \neq \tilde{\mathbf{p}}^i, \text{ and } \mathbf{p}^j = \tilde{\mathbf{p}}^j \text{ for all } j \neq i. \\ 1 - \displaystyle\sum_{\tilde{\mathbf{p}}^i \neq \mathbf{p}^i} \dfrac{1}{n(k-1)} \cdot \rho_s(\pi_i, \tilde{\pi}_i) & \text{if } P = \tilde{P} \\ 0 & \text{otherwise.} \end{cases} \tag{27}
$$

The transition probabilities according to Eq. (27) indicate that within each time step, at most one player can change her strategy. This player is chosen randomly (a particular player i is chosen with probability $1/n$), and then given the chance to pick a new strategy (each remaining strategy has the same probability $1/(k-1)$). This strategy is then adopted with probability $\rho_s(\pi_i, \tilde{\pi}_i)$. Again, we can collect all transition probabilities in a $k^n \times k^n$ matrix $W = (w_{P,\tilde{P}})$. If the initial state of the population is given by the vector $\mathbf{y}(0) = (y_P(0))$, then the expected strategy distribution at time t is $\mathbf{y}(t) = \mathbf{y}(0) \cdot W^t$. Because the evolutionary process is ergodic for all finite selection strengths s, the distribution $\mathbf{y}(t)$ converges to a unique invariant distribution $\mathbf{y} = (y_P)$ that solves $\mathbf{y} = \mathbf{y} \cdot W$. Its entries y_P reflect how often we are to observe the population in state $P = (\mathbf{p}^1, \ldots, \mathbf{p}^n)$ over an evolutionary timescale.

For the evolutionary results depicted in our Extended Data Figures, we have employed this Markov chain approach for groups of size $n = 2$. In that case, each player can choose among $k = 2^5 = 32$ pure memory-1 strategies (when $\delta < 1$), or among $k = 2^4 = 16$ strategies (when $\delta = 1$, such that the cooperation probability in the first round can be neglected). **Extended Data Fig. 1** is based on the five scenarios considered in the behavioral experiment. It compares realized cooperation trajectories of introspection dynamics as defined in Section 4.2 to the expected trajectories defined in this section. **Extended Data Fig. 2** depicts the invariant distribution $\mathbf{y}$ for the same five scenarios. In three cases, we observe that players predominantly settle at a state where both players employ Win-Stay Lose-Shift (WSLS, see next subsection). To calculate the invariant distribution $\mathbf{y}$, we have used MATLAB to solve the homogeneous linear system $(I_{256} - W')\mathbf{y} = 0$. Here, I_{256} is the $2^4 \cdot 2^4 = 256$-dimensional identity matrix, and W' refers to W transposed. **Extended Data Fig. 3** shows the players' average contributions and how often they employ WSLS, again computed from the invariant distribution $\mathbf{y}$ of the process. **Extended Data Fig. 4** illustrates how the players' average payoffs according to the invariant distribution $\mathbf{y}$ change as we vary the expected number of rounds $1/(1-\delta)$, the selection strength s, or the error rate ε. Finally, in **Extended Data Fig. 5** we present evolutionary results when players are engaged in a public goods game with asymmetric benefits, with payoffs given by Eq. (13). The MATLAB implementation to generate the respective data is provided in the **Appendix**.

4.4 Analysis of Win-Stay Lose-Shift

Equilibrium conditions for Win-Stay Lose-Shift. For our evolutionary results, we have assumed that the players' decisions are subject to noise. Under such conditions, a homogeneous group of Grim players may no longer be able to maintain full cooperation: if one of the Grim players fails to cooperate due to an error, the whole group switches to defection for the remainder of the game. Instead, the evolutionary results depicted in **Extended Data Fig. 1 – Extended Data Fig. 5** suggest that for high continuation probabilities, full cooperation only emerges as the strategy Win-Stay Lose-Shift (WSLS) becomes predominant.

While previous studies have explored this strategy for symmetric social dilemmas with discrete actions[49–52], we can extend the definition of WSLS for general n-player public goods games with continuous contribution levels as follows. A WSLS player contributes her full endowment in the first round, or if all players with a positive endowment have contributed the same share of their endowment in the previous round. Otherwise, if people have differed in their relative contributions in the last round, a WSLS player contributes nothing. For a two-player game with two discrete contribution levels, this definition reduces to the usual definition of WSLS, $\mathbf{p} = (p_0;\ p_{CC}, p_{CD}, p_{DC}, p_{DD}) = (1;\ 1, 0, 0, 1)$.

In contrast to Grim, a population of WSLS players is robust with respect to noise. If one of the players deviates from full cooperation due to an error, all players collectively defect the next round, and then revert to mutual cooperation in all subsequent rounds.

The following Proposition characterizes when WSLS is an equilibrium. The proofs of the results in this section are again given in the **Appendix**.

Proposition 9. *Consider a public good game with payoff function u, continuation probability $\delta < 1$, and endowment distribution $\mathbf{e}$. Then the strategy profile where all players apply WSLS is a subgame perfect equilibrium if and only if*

$$\delta\Big(u_i(\mathbf{e}, \mathbf{1}) - u_i(\mathbf{e}, \mathbf{0})\Big) \geq \Big(u_i(\mathbf{e}, \mathbf{1}_{-i}) - u_i(\mathbf{e}, \mathbf{1})\Big) \quad \textit{for all players } i \textit{ with } e_i > 0. \tag{28}$$

We note that Proposition 9 does not make any restrictions on the possible deviation strategies. When condition (28) holds, there is no incentive to deviate from WSLS even if players have access to higher-memory strategies, or when they choose among continuous contribution levels.

When we compare the two Propositions 1 and 9, we note that the critical threshold δ for WSLS to be an equilibrium is strictly larger than the respective threshold for Grim. As a consequence, the set of endowment distributions for which WSLS is an equilibrium, referred to as $E_u^W(\delta)$, is a strict subset of the set $E_u(\delta)$ of endowments where Grim is an equilibrium. This relationship is also illustrated in **Extended Data Fig. 3**, where the region $E_u^W(\delta)$ (indicated by dashed lines) is strictly contained in the region $E_u(\delta)$ (indicated by dotted lines). As shown in that figure, high cooperation rates only evolve in the parameter region where WSLS is a subgame perfect equilibrium. The following Proposition shows that the sets $E_u^W(\delta)$ and $E_u(\delta)$ have similar qualitative properties.

Proposition 10.

1. *Consider a public goods game with payoff function u, and let δ, δ' be two continuation probabilities with $\delta < \delta'$. Then $E_u^W(\delta) \subset E_u^W(\delta')$.*

2. *If $E_u^W(\delta) \neq \emptyset$ for some public goods game with linear and symmetric payoff function u and for some continuation probability δ, then $(1/n, \ldots, 1/n) \in E_u^W(\delta)$.*

3. *There are public good games with an asymmetric payoff function u such that $E_u^W(\delta) \neq \emptyset$ for some continuation probability δ, but $(1/n, \ldots, 1/n) \notin E_u^W(\delta)$.*

4. *Consider a pairwise game with linear payoff function (1). If $E_u^W(\delta) \neq \emptyset$ for some given continuation probability δ, then $\mathbf{e}^* \in E_u^W(\delta)$, where $\mathbf{e}^*$ is the maximally cooperative endowment distribution as characterized by Eq. (21).*

To the extent that the evolution of cooperation is linked to the performance of WSLS, the above results complement the static results from the previous section. Cooperation should more readily evolve as the continuation probability δ increases; in symmetric and linear games, cooperation is most likely to evolve under an equal endowment distribution; and in asymmetric games, an uneven distribution of endowments is more favorable to the evolution of cooperation. These predictions based on the stability of WSLS are further supported by our numerical results in **Fig. 2**, **Fig. 3**, and **Extended Data Fig. 1** – **Extended Data Fig. 5**.

Evolutionary stability of Win-Stay Lose-Shift. The conditions for subgame perfection in Eq. (28) guarantee that when all players adopt WSLS, no player can gain a higher payoff by deviating. However, they do not imply that WSLS is evolutionarily stable in the sense of Maynard Smith[53]. After all, neutral deviations that leave the deviating player with exactly the same payoff may still be possible.

There is by now a rather extensive literature on whether evolutionarily stable strategies for repeated games exist at all[54–63]. For the repeated prisoner's dilemma without errors, Selten and Hammerstein have noted that Tit-for-Tat is not evolutionarily stable because it can be neutrally invaded by ALLC[54]. Boyd and Lorberbaum have generalized this result to show that in the absence of errors, no pure strategy is evolutionarily stable[55]. In particular, they show that once a first neutral mutant gains a sizable proportion of the population, a second mutant strategy can be selectively favored to invade. Similar results apply to mixtures of pure strategies[56] and to stochastic strategies[57]. Taking a dynamic perspective, van Veelen and Garcia have shown in a series of papers that cooperation rates in evolving populations tend to oscillate[59–61]. For any resident population, there always exist neutral 'stepping stone' mutations out of equilibrium. Over time, cooperation comes and goes.

However, these results do not mean that evolutionary stability in repeated games is generally impossible. Under the realistic assumption that players sometimes commit errors, Boyd has shown that a strategy is evolutionarily stable if it is a 'strong perfect equilibrium'[62]. That is, in a game with discounted payoffs, a strategy is evolutionarily stable if it is a strict best response to itself after any possible history. Intuitively, errors guarantee that any possible history is reached with positive probability. This general reachability in turn prohibits neutral mutants to exist that only differ from the resident off the

equilibrium path. Based on this insight, we can reformulate Proposition 9 to gain the following result on the evolutionary stability of WSLS.

Proposition 11. *Consider a public good game with payoff function u, continuation probability $\delta < 1$, and endowment distribution $\mathbf{e}$ with $e_i > 0$ for all players i. Moreover, suppose the inequality (28) is strict,*

$$\delta\Big(u_i(\mathbf{e}, \mathbf{1}) - u_i(\mathbf{e}, \mathbf{0})\Big) > \Big(u_i(\mathbf{e}, \mathbf{1}_{-i}) - u_i(\mathbf{e}, \mathbf{1})\Big) \quad \textit{for all players } i. \tag{29}$$

Then there exists a threshold $\varepsilon_0 > 0$ such that WSLS is evolutionarily stable for all error rates $0 < \varepsilon < \varepsilon_0$.

When the conditions of Proposition 11 are met, any deviation in a group of WSLS players leads to a strictly lower expected payoff for the deviating player. As a consequence, any such deviation is disfavored by the evolutionary process.

Two remarks are in order. First, we stress that Proposition 11 again does not restrict the possible mutant strategies. Even if players are permitted to react to arbitrarily long histories of past play, any deviation from WSLS leaves them with a strictly lower payoff. Second, we note that for any finite selection strength, the evolutionary process we have considered is ergodic. That is, no matter which strategy profile the group currently adopts, there is always a chance that one of the players decides to switch his strategy, even if the new strategy yields a strictly lower payoff. The process is guaranteed to leave any strategy profile eventually. Realized cooperation trajectories will thus exhibit oscillations, even if WSLS is evolutionarily stable (see **Extended Data Fig. 1** for some typical examples of such trajectories).

Nevertheless, the conditions in Proposition 11 are important, as they affect how often players cooperate on average over an evolutionary timescale. Once WSLS is evolutionarily stable, a group of WSLS players can only be left due to a rare chance event. Depending on the strength of selection, a considerable number of elementary introspection events may be necessary until one player deviates.

4.5 Evolution among players with finite state automata

So far we have explored the evolution of cooperation when players can choose among all memory-1 strategies. Memory-1 strategies have been routinely used in previous work as they are convenient to analyze[17–22]. However, the emergence of reciprocity in evolutionary simulations can be sensitive to the considered strategy space[50,64,65] and to the applied mutation scheme[66,67]. It is thus natural to ask to which extent our evolutionary results depend on the assumption of memory-1 players. To explore this issue in more detail, we consider the evolutionary dynamics when players have access to a vastly larger strategy set.

Strategies represented by finite state automata. In the following, we adapt the framework of van Veelen and Garcia[59–61] to the asymmetric games studied herein. For simplicity, we describe the framework for pairwise games only, and we assume players use pure strategies without errors.

As before, players can choose among a finite set of contribution levels $X = \{\hat{x}_1, \dots, \hat{x}_m\}$. To make their decisions, players use finite state automata $\mathcal{A} = (\Omega, \lambda, \mu)$. The first component $\Omega = \{\omega_1, \dots, \omega_k\}$ is the set of possible states of the automaton. The second component $\lambda : \Omega \to X$ determines how much the player contributes in each state. Finally, the third component $\mu : \Omega \times X \to \Omega$ is the transition function. Depending on the automaton's present state and on the co-player's contribution, it determines the player's state in the next round. Without loss of generality, we assume that the finite state automaton is initially in the first state ω_1.

The memory-1 strategies considered in the previous sections can all be represented as finite state automata. As an example, let us assume that players can either contribute their full endowment or nothing, $X = \{0, 1\}$. Again, we identify these two actions with cooperation (C) and defection (D), respectively. Suppose a player uses the memory-1 strategy $\mathbf{p} = (p_0; \ p_{CC}, p_{CD}, p_{DC}, p_{DD})$, and for simplicity assume that initially the player cooperates, $p_0 = 1$. Then $\mathbf{p}$ is equivalent to a finite state automaton $\mathcal{A}_{\mathbf{p}} = (\Omega_{\mathbf{p}}, \lambda_{\mathbf{p}}, \mu_{\mathbf{p}})$ with elements defined by

$$\Omega_{\mathbf{p}} = (\omega_1, \omega_2), \quad \lambda_{\mathbf{p}}(\omega_1) = C \text{ and } \lambda_{\mathbf{p}}(\omega_2) = D, \quad \mu_{\mathbf{p}}(\omega, x) = \begin{cases} \omega_1 & \text{if } p_{\lambda_{\mathbf{p}}(\omega), x} = 1 \\ \omega_2 & \text{if } p_{\lambda_{\mathbf{p}}(\omega), x} = 0. \end{cases} \tag{30}$$

We illustrate this representation of memory-1 strategies as finite state automata in **Extended Data Fig. 6a**. There, we depict three particular examples, AllD, Tit-for-Tat, and Win-Stay Lose-Shift.

While every memory-1 strategy can be represented as a finite state automaton, the set of strategies representable by finite state automata is considerably richer. As an example of a strategy that is not memory-1, **Extended Data Fig. 6a** provides a finite-state representation of Tit for two Tats[68]. More generally, it can be shown that in fact all possible pure strategies for the repeated prisoner's dilemma can be approximated arbitrarily closely by finite state automata[60].

To calculate the players' payoffs, suppose player 1 uses an automaton with k_1 states and player 2 uses an automaton with k_2 states. Because strategies are pure and players do not commit errors, the game dynamics is deterministic. In particular, it takes at most $k_1 \cdot k_2$ rounds until the two players enter an infinite loop. From that moment on, their actions repeat themselves indefinitely. For the simulations shown in **Extended Data Fig. 6c**, we assume no discounting, and the players' payoffs are thus simply defined as their average payoff per round over all rounds of the loop. Similar results can be obtained in discounted games. In the **Appendix**, we provide the MATLAB implementation that we have used.

Introspection dynamics for players with finite state automata. To model the evolutionary dynamics, we assume that initially both players use the AllD automaton depicted in **Extended Data Fig. 6a**. In each time step, one of the players is then chosen at random to experiment with a mutation $\tilde{\mathcal{A}}$ of his present automaton $\mathcal{A}$. Following van Veelen and Garcia[59,60], we allow for four different types of mutations (see also **Extended Data Fig. 6b**),

(*i*) Mutation in the state transitions ('Change an arrow'). When this mutation occurs, one input (ω, x) of the transition function is randomly chosen. Then its output $\mu(\omega, x)$ is changed randomly, such

that all alternative states in Ω have the same probability $1/(k-1)$ to become the new output.

(*ii*) <u>Mutation in the chosen contribution level ('Change an action').</u> When this mutation occurs, one state $\omega \in \Omega$ is randomly chosen. Then the chosen contribution level $\lambda(\omega)$ in that state is changed randomly, such that all alternative contribution levels have the same probability $1/(m-1)$ to be chosen.

(*iii*) <u>Mutation in the number of states ('Remove a state').</u> When this mutation occurs, one state $\omega \in \Omega$ is randomly chosen. This state is then removed from Ω. All transitions that lead into this state are randomly and independently redirected. All remaining states have the same probability to be the new target of these redirected transitions.

(*iv*) <u>Mutation in the number of states ('Add a state').</u> When this mutation occurs, a new state ω' is added to the set Ω. For each pair (ω', x), the outgoing transitions $\mu(\omega', x)$ are determined randomly, such that all states in Ω have the same probability to be the target (including the new state). In addition, one of the existing transitions among the previous states is randomly chosen, and randomly redirected towards the new state. That is, a pair (ω, x) with $\omega \neq \omega'$ is randomly chosen, and the value of $\mu(\omega, x)$ is changed to ω'.

As in the case of memory-1 strategies, we assume that the player then compares his previous payoff π_i with the payoff $\tilde{\pi}_i$ he could have got when using the mutated automaton $\tilde{\mathcal{A}}$. The player switches to the mutated automaton with probability $\rho_s(\pi_i, \tilde{\pi}_i)$, as defined in Eq. (26). This elementary updating process is then again iterated for T time steps. This yields two sequences $\left(\mathcal{A}_1(0), \ldots, \mathcal{A}_1(T)\right)$ and $\left(\mathcal{A}_2(0), \ldots, \mathcal{A}_2(T)\right)$, where $\mathcal{A}_i(t)$ is the automaton used by player i after t time steps. Based on this sequence, we calculate the average cooperation rate over the entire simulation.

We note that the above mutation scheme is capable of producing every possible finite state automaton over the set of possible contributions X. Like van Veelen and Garcia[59,60], we assume for the simulations that the four types of mutations mentioned occur with probabilities 17.5%, 17.5%, 35%, and 30%, respectively. We obtain similar results when different probabilities are used (however, for computational reasons, it is useful to assume that the removal of a state is more likely than adding a new state).

Simulation results. To compare the evolutionary dynamics among players with finite state automata with our previous results on memory-1 strategies, we have re-run the simulations of **Fig. 3** in the main text. That is, we consider two scenarios, depending on whether the players' productivities are equal or unequal ($r_1 = r_2 = 1.6$ compared to $r_1 = 1.9 > 1.3 = r_2$). In each scenario, we run separate simulations for different endowment distributions, $e_1 \in \{0, 0.1, \ldots, 0.9, 1\}$. The results are shown in **Extended Data Fig. 6c**.

Compared to **Fig. 3**, we first observe that the simulations based on finite state automata yield somewhat lower cooperation rates overall. This reduced level of cooperation seems to be partly driven by the differences in the mutation process. In the memory-1 space, we have assumed that every possible

strategy is equally likely to be the next mutant. In particular, an AllD player may switch to Tit for Tat within one time step. In contrast, when players use finite state automata, an AllD player can only arrive at Tit-for-Tat through a sequence of consecutive elementary mutations. This sequence is easily interrupted once a mutation leads to the removal of a state. As a consequence, we observe that players with finite state automata tend to defect unconditionally more often than memory-1 players.

However, all qualitative results reported before remain unchanged. When players are equally productive, we observe that cooperation is most abundant when players obtain equal endowments. Conversely, when players have different productivities, cooperation rates reach a maximum when the more productive player obtains more of the endowment. Again, the position of this maximum is reasonably approximated by the maximally cooperative endowment distribution as computed in Eq. (22).

Overall, the results in **Fig. 3** and **Extended Data Fig. 6** suggest that our evolutionary results on the effect of endowment inequality are not sensitive to the considered strategy space. Provided that the strategy space is kept constant, equal endowments promote cooperation when players are equally productive. Unequal endowments prove beneficial when players differ in their productivity.

5 Experimental analysis

5.1 Methods and procedures

Recruitment and sample size. To explore the validity of our theoretical results, we have recruited $N = 436$ participants on Amazon Mechanical Turk (AMT). AMT is an online labor market which is routinely used for psychological and economic studies[69–71]. Participants on AMT receive a base payment for participation: the base pay was \$1.00 in our experiment. In addition, participants may receive further payments, known as a "bonus payment", based on their performance in the game. To earn a bonus, participants could earn tokens during the game which were converted to US dollars at the end of the game at a rate of 800 tokens = \$1.00; the average bonus participants earned was \$1.70.

For our experimental study, participants were recruited to interact with each other in pairs. We implemented the online experiment with SoPHIE, an online platform that allows for real-time interaction between AMT participants[72–74]. The sample size was determined in advance based on similar past research[74]. We recruited at least 80 participants per condition because we expected that the effect size for our key predictions would be moderate and that the dropout rate would be about 30%, as previously observed in similar interactive online experiments[73,74]. In the end, dropouts were lower than expected (about 7%; for a detailed dropout analysis, see Section 5.3).

The experiment was approved by the Harvard University Institutional Review Board (Protocol Number 14306).

Instructions and comprehension. Participants entered the study, read the instructions and had to pass several comprehension questions before they could continue (the instructions as well as the questions are provided in the **Appendix**). If they failed one of the comprehension questions, the relevant instructions

were shown again so that participants could try again. Only when all questions have been answered successfully were participants eligible to continue.

After reading the instructions and passing the comprehension questions, participants entered a waiting room. Once at least two participants had arrived, they were paired for the duration of the experiment. If no other participant arrived within 10 minutes, unpaired participants were redirected to an exit screen, and thanked for their attention and patience. They received the respective $1.00 base payment for participating in the study but no further bonus payment.

Parameters and decisions each round. We implemented five different treatments based on the linear public good games (7) and (10). The treatments differed in how the initial endowment was distributed among the players, and in how productive the players' contributions were (i.e., by which factor a player's absolute contributions were multiplied). Specifically, we have used the following parameters (see also **Fig. 4a**):

Treatment	Endowment e_i		Productivity r_i	
	Player 1	Player 2	Player 1	Player 2
(1) Full equality	50	50	1.6	1.6
(2) Endowment inequality	75	25	1.6	1.6
(3) Productivity inequality	50	50	1.9	1.3
(4) Aligned inequality	75	25	1.9	1.3
(5) Misaligned inequality	25	75	1.9	1.3

The parameter values for the treatments $(1) - (3)$ were chosen such that (i) the total size of the endowment is constant ($e_1 + e_2 = 100$), and (ii) the maximum achievable group payoff is the same across conditions. The parameter values for the treatments (4) and (5) then follow from symmetry considerations. We note that all parameter values are in line with the four properties **(C)** – **(OC)**, as introduced in Section 1.

Every pair of participants interacted in at least 20 rounds of the public goods game. To avoid end-game effects, players knew that in each round after the 20th round, there was a 50% probability that there would be another round. From this distribution, the average number of rounds was 21.0 with a range between 20 and 29. Taking into account that due to dropouts some groups did not complete the game, the average number of rounds actually played was 20.3 with a range between 1 and 29.

In each round, both participants had to decide simultaneously how much to contribute to the public good. Participants could contribute any integer value between 0 and e_i, the participant's respective endowment. After both participants made their decisions, they received all relevant information on the outcome of this round: each participants' absolute contributions, their productivity factor, the total public good provided by each participant, and the returns from the public good per participant. Participants then continued to the next round of the game until the end of the game. When the game had ended, participants were thanked for their participation, and they received their payment.

Statistical methods. For all our main experimental results, depicted in **Fig. 4** and **Extended Data Figs. 7–10**, we used non-parametric tests (Mann-Whitney test for between-treatment comparisons, and Wilcoxon signed-rank test for within-treatment comparisons). For these tests, we considered groups of two players as our statistical unit. We used two-tailed tests throughout. For between-treatment comparisons, we restrict ourselves to comparisons where the players' productivity is kept fixed but their endowments may vary. That is, we compare the two treatments (1) and (2), and we separately compare the treatments (3), (4), and (5). In this way our results indicate, for given productivities of the players, how a redistribution of endowments would affect the players' contributions to the public good.

When we report results on the players' contributions, we consider their *relative* contributions (i.e., which fraction of their endowment the players contribute, corresponding to the variable x in our theoretical model). As a consequence, we note that even when we find that a high-endowment player gives relatively less than a low-endowment player ($x_1 < x_2$), the high-endowment player's absolute contributions may still be higher (when $e_1 x_1 > e_2 x_2$).

For our analysis, we only consider the first twenty rounds of the game because this was the minimum number all groups had in common. Moreover, for the results in the main text and in the following section, we have excluded groups in which at least one participant dropped out half-way through the game. However, in Section 5.3 we repeat our analysis using imputed values for the dropout groups instead. All our main results remain unchanged. Unless noted otherwise, all statistical tests have been performed with Stata SE 15.1.

5.2 Results

Average payoffs and contributions across treatments. In **Fig. 4**, we have depicted the players' average contributions across the five treatments, as well as the resulting surplus (i.e., the difference between the players' payoffs and their initial endowments, relative to the initial endowments).

When players are equally productive, we find that their total contributions make up 72.2% of their endowment in the full equality treatment, compared to 49.0% under endowment inequality (Mann-Whitney test, $Z = 2.813$, $p = 0.005$). As a direct consequence, the two treatments also differ in the generated surplus, which is 43.3% under full equality but only 29.4% under endowment inequality (Mann-Whitney test, $Z = 2.813$, $p = 0.005$).

For the case where players have unequal productivities, we find no significant differences in the players' contributions between productivity inequality and aligned inequality: players contribute 72.8% of their endowment under productivity inequality and 73.5% under aligned inequality (Mann-Whitney test, $Z = 0.844$, $p = 0.399$). However, both of these treatments differ significantly from the treatment with misaligned inequality, where players only contribute 45.5% of their endowment (Mann-Whitney tests, treatments (3) vs. (5): $Z = 3.818$, $p < 0.001$; (4) vs. (5): $Z = 3.825$, $p < 0.001$). With regard to the generated surplus all three treatments differ from one another. As predicted by our theoretical analysis, aligned inequality yields the highest surplus (the respective values are (3): 43.8%, (4): 55.1%, (5): 23.1%; Mann-Whitney tests, (3) vs. (4): $Z = 3.357$, $p < 0.001$; (3) vs. (5): $Z = 4.592$, $p < 0.001$;

(4) vs. (5) $Z=4.802$, $p<0.001$).

Comparing the contributions of the two players. In **Extended Data Fig. 7** we show the contributions and payoffs for each treatment, separately for the two players.

As one may expect, in the full equality treatment there are no differences in the players' contributions and payoffs. The average contribution of players 1 and 2 are 72.3% and 72.2%, respectively, and their payoffs are 71.63 tokens and 71.69 tokens per round, respectively (Wilcoxon test, $Z=0.784$, $p=0.433$).

Under endowment inequality, the relative contributions of the high-endowment player 1 fall below the contributions of player 2, their average contributions being 47.2% and 54.4%, respectively (Wilcoxon text, $Z=2.194$, $p=0.028$). As a consequence, the high-endowment player keeps more of their own units and thus earns a substantially higher payoff of 78.8 tokens, compared to the 50.6 tokens for the low-endowment player (Wilcoxon test, $Z=5.098$, $p<0.001$).

Under productivity inequality, we find that the relative contributions of the two players are similar, with the more productive player 1 contributing 73.4% of her endowment on average, and the less productive player 2 contributing 72.2% (Wilcoxon test, $Z=0.380$, $p=0.704$). As a result, also their payoffs are comparable, 71.6 tokens and 72.2 tokens, respectively (Wilcoxon test, $Z=0.366$, $p=0.714$).

We find a similar agreement in the players' contributions under aligned inequality. Here the high-endowment and high-productivity player 1 contributed 73.4% of her endowment, whereas the low-endowment and low-productivity player 2 contributed 73.8% (Wilcoxon test, $Z=0.705$, $p=0.481$). However, since the players are not fully cooperative and player 1's initial endowment is higher, player 1 obtains a higher payoff on average, with 84.2 tokens compared to 70.8 tokens (Wilcoxon test, $Z=3.239$, $p=0.001$).

Finally, under misaligned inequality, the players differ in both their contributions and their eventual payoffs. Whereas player 1 contributes 62.9% of her endowment and earns a payoff of 43.5 tokens, player 2 contributes 39.7% of his endowment and earns 79.5 tokens (Wilcoxon tests, contributions: $Z=4.234$, $p<0.001$; payoffs: $Z=5.464$, $p<0.001$).

Dynamics of cooperation over time. In **Extended Data Fig. 8** we show how much the two players in each treatment have contributed over time. To test whether there is any significant time trend within a treatment, we have compared the players' total contributions during the first five rounds with their total contributions during the last five rounds. For four of the five treatments we find no significant trends over time (Wilcoxon tests, $p>0.1$ for all four of these treatments); only under endowment inequality, we find that overall contributions slightly decrease, from 54.0% during the first five rounds to 43.9% during the last 5 rounds (Wilcoxon test, $Z=2.578$, $p=0.010$).

Reciprocal behavior across the five treatments. To gain some insights into the players' strategies, we have analyzed to which extent their behaviors are guided by reciprocity considerations. To this end, we say that a player's relative contribution is reciprocal if it matches or exceeds the co-player's relative

contribution in the previous round. That is, if one player contributes 60% of her endowment in one round, a reciprocal co-player would contribute at least 60% of his endowment in the following round. As illustrated in **Extended Data Fig. 10**, a majority of the players' contribution decisions can be classified as reciprocal, such that the frequency of reciprocal behavior is typically between 70% and 90%.

The only exception is the treatment with misaligned inequality. Here, the low-endowment high-productivity player 1 reciprocates in 85.9% of all rounds, whereas the high-endowment low-productivity player 2 only reciprocates in 44.9% of all cases (Wilcoxon test, $Z = 4.252$, $p < 0.001$). To explore this mismatch in the players' reciprocity in more detail, **Extended Data Fig. 9** shows the players' joint contribution decisions. For misaligned inequality, we observe that in a substantial fraction of rounds, the two players coordinate on a state where they both give the same absolute amount to the public good (for example, in 12.4% of the cases, both players contribute 25 tokens, which is the maximum contribution for player 1, but only a third of player 2's endowment). As a result, player 1's relative contributions systematically exceed the relative contributions of player 2.

5.3   Robustness of the statistical results

As is usual in online experiments, some groups did not complete the study because players decided to drop out. These instances could potentially confound our analysis if there is a correlation between the game situation a participant experiences and his decision to leave the experiment. To quantify the effect of dropouts, we have performed two separate analyses. First, we have explored whether there are some systematic patterns in the groups that dropped out. Second, we have repeated all our statistical tests using imputed values for the groups that did not complete the study. Both analyses suggest that our statistical results are robust with respect to attrition.

Dropout analysis. Here, we investigate whether selection bias may have occurred on the basis of exogenous or endogenous variation (e.g. depending on the assigned treatment, or on the experienced contribution behavior during the game). In total, fifteen out of 218 groups (6.9%) did not complete the game because at least one of the players dropped out. The following table shows how many groups completed the experiment in each treatment.

Treatment	Completed game		Total
	No	Yes	
(1) Full equality	2	42	44
(2) Endowment inequality	1	42	43
(3) Productivity inequality	3	40	43
(4) Aligned inequality	5	39	44
(5) Misaligned inequality	4	40	44
Total	15	203	218

Most dropouts occurred in the treatment with aligned inequality (4), whereas the fewest groups dropped out in the endowment inequality treatment (1). When we use the data of this table to test whether dropouts depend on the treatment, we find no significant association (Fisher's exact test, $p = 0.531$).

Next we have explored whether dropouts are associated with other variables, such as the group's cooperation dynamics. To this end, we have explored several logistic regression models. Because dropouts are comparably rare, standard logistic regression can fail due to the problem of "separation"[75]. In that case, the estimate for the regression parameter diverges and the maximum likelihood estimate does not exist. We observe this divergence for our endowment inequality treatment, for which there is only a single group that dropped out. To address this problem, we use Firth's logistic regression[76] instead, with standard errors clustered at the group level. We have verified that in those cases in which standard logistic regression produces finite parameter estimates, the results are qualitatively similar.

We have first explored whether less cooperative groups are more likely to drop out. As the independent variable we use the average sum of contributions across all rounds for each group. This variable takes values between 0 (if neither player ever contributed anything) and 100 (if both players always contributed everything). As the dependent variable we use an indicator variable with values 0 or 1 (dropped out or completed the game, respectively). When we pool across all conditions, there is no significant relationship between the dropout status of a group and the average contributions in a group ($p = 0.309$).

Next, we add a variable for the condition assignment to this logistic regression (1 = full equality, 2 = endowment inequality, 3 = productivity inequality, 4 = aligned inequality and 5 = misaligned inequality). As before, average contributions do not predict dropout ($p = 0.409$). Furthermore, none of the condition indicators are significant (p-value of (2) endowment inequality dummy: $p = 0.754$; (3) productivity inequality: $p = 0.655$; (4) aligned inequality: $p = 0.269$; (5) misaligned inequality: $p = 0.517$). These results suggest that there is no significant relationship across conditions between average contributions and whether or not the group completes the game.

We also consider each treatment separately. We use Firth's logistic regression with the same independent and dependent variables at the group level, but now for each treatment. The sum of contributions is not a significant predictor of group dropout in the full equality ($p = 0.579$), endowment inequality ($p = 0.313$), productivity inequality ($p = 0.736$), aligned inequality ($p = 0.271$) or misaligned inequality ($p = 0.319$) conditions.

Finally, we explore whether treatment assignment and contribution dynamics jointly affect dropouts. To this end, we pool across all conditions. As the independent variables, we use

(*i*) the sum of contributions (ranging between 0 and 100, abbreviated as SOC below),

(*ii*) the five condition indicators (1 = full equality, 2 = endowment inequality, 3 = productivity inequality, 4 = aligned inequality, 5 = misaligned inequality; full equality is also the baseline group for comparisons) and

(*iii*) the interaction term between these two variables.

The dependent variable is 0 if the group dropped out or 1 if the group completed the game. We find

that none of the interaction terms are significant predictors of dropout (SOC interacted with endowment inequality indicator: $p=0.646$; with productivity inequality indicator: $p=0.815$; with aligned inequality indicator: $p=0.330$; with misaligned inequality indicator: $p=0.881$). A post-hoc comparison of the interaction coefficients also does not reveal any significant differences (comparisons of interaction terms: SOC $\times$ (2) vs. SOC $\times$ (3): $p=0.487$; SOC $\times$ (2) vs. SOC $\times$ (4): $p=0.185$; SOC $\times$ (2) vs. SOC $\times$ (5): $p=0.705$; SOC $\times$ (3) vs. SOC $\times$ (4): $p=0.386$; SOC $\times$ (3) vs. SOC $\times$ (5): $p=0.652$; SOC $\times$ (4) vs. SOC $\times$ (5): $p=0.152$).

Overall, we find no indication that there is a systematic bias in the players' dropout decisions.

Robustness analysis with imputed missing values. For the experimental results reported in Section 5.2 we have only considered groups that have completed the game. In the following, we provide a complementary analysis in which we include missing data using a multiple imputation scheme.

Specifically, we apply the iterative, non-parametric imputation method introduced by Stekhoven and Bühlmann[77]. Their 'missForest' imputation method uses random forests (RF), which is a supervised machine learning method based on an iterative decision tree process. In our case, the generated trees split at nodes of variables that predict contributions on the basis of the existing data and makes predictions about the missing values. As is common with RF procedures, the final dataset of predicted values is then tested and validated against a hold-out sample.

As predictors we use treatment assignment, round, player type, endowment, productivity and non-missing contributions. Using these variables, we predict out-of-sample the missing contributions of dropout groups. The parameters for the missForest function in R were as follows. We kept the standard number of trees supplied to the RF method (100), because our data is not high-dimensional. Decreasing or increasing the number of trees did not change the results. We set the maximum iteration parameter for the imputation procedure to 10 such that after 10 iterations, the procedure is forced to a hard-stop if it has not met its stopping criteria. The stopping criteria are such that after each iteration, the previous and the current imputed data matrix are compared. If the difference between them goes up for the first time, the most recent imputed dataset is returned. In practice, the stopping criteria were met after 4 iterations and the hard-stop was not enforced. The missForest function returns the normalized root mean squared error (NRMSE) to assess how the final imputed dataset fared in the out-of-sample prediction; the NRMSE was 0.000016 in our computations, suggesting a low error rate.

With the imputed dataset, we repeat the same analysis as in the previous Section 5.2. The results are summarized in **Supplementary Table 1** (for the two treatments in which players have equal productivities) and in **Supplementary Table 2** (for the three treatments in which players have different productivities). In each case, the left column ('main sample') displays our previous results based on only those groups who completed the study. The right column ('Imputed sample') shows the new results using the data of all groups, based on the imputation method described above. As before, we use non-parametric Mann-Whitney tests for between-treatment comparisons, and Wilcoxon signed-rank tests for within-treatment comparisons (two-tailed).

31

	Main sample	**Imputed sample**
Contributions	**(1)** 72.2% vs. **(2)** 49.0%, $Z = 2.813$**	**(1)** 73.4% vs. **(2)** 50.2%, $Z = 2.847$**
Surplus	**(1)** 43.3 vs. **(2)** 29.4, $Z = 2.813$**	**(1)** 44.0 vs. **(2)** 30.1, $Z = 2.847$**
Player 1 vs. Player 2 Contributions	**(1)** 72.3% vs. 72.2%, $Z = 0.765$ **(2)** 47.2% vs. 54.4%, $Z = 2.194$*	**(1)** 73.5% vs. 73.3%, $Z = 0.524$ **(2)** 48.5% vs. 55.4%, $Z = 2.234$*
Player 1 vs. Player 2 Payoffs	**(1)** 71.6 vs. 71.7, $Z = 0.784$ **(2)** 78.8 vs. 50.6, $Z = 5.098$***	**(1)** 72.0 vs. 72.1, $Z = 0.548$ **(2)** 78.8 vs. 51.3, $Z = 5.159$***
First versus last **five rounds**	**(1)** 71.1% vs. 72.8%, $Z = 0.350$ **(2)** 54.0% vs. 43.9%, $Z = 2.578$**	**(1)** 72.3% vs. 73.9%, $Z = 0.367$ **(2)** 55.0% vs. 45.1%, $Z = 2.653$**
Player 1 vs. Player 2 Reciprocity	**(1)** 90.0% vs. 88.5%, $Z = 0.804$ **(2)** 71.6% vs. 80.1%, $Z = 1.449$	**(1)** 90.3% vs. 88.5%, $Z = 0.552$ **(2)** 72.1% vs. 80.5%, $Z = 1.533$

Supplementary Table 1: Results for main and imputed samples under equal productivites. Bold numbers in brackets refer to the respective treatment, **(1)** Full equality, **(2)** Unequal endowments. The number of data points is $N_1 = N_2 = 42$ (left column) and $N_1 = 44$, $N_2 = 43$ (right column), respectively.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, using either a Mann-Whitney test (for between-treatment comparisons) or a Wilcoxon signed-rank test (for within-treatment comparisons). All tests are two-tailed.

	Main sample	**Imputed sample**
Contributions	**(3)** 72.8% vs. **(4)** 73.5%, $Z = 0.844$ **(3)** 72.8% vs. **(5)** 45.5%, $Z = 3.818$*** **(4)** 73.5% vs. **(5)** 45.5%, $Z = 3.825$***	**(3)** 73.5% vs. **(4)** 72.4%, $Z = 0.420$ **(3)** 73.5% vs. **(5)** 44.5%, $Z = 4.226$*** **(4)** 72.4% vs. **(5)** 44.5%, $Z = 4.114$***
Surplus	**(3)** 43.8 vs. **(4)** 55.1, $Z = 3.357$*** **(3)** 43.8 vs. **(5)** 23.1, $Z = 4.592$*** **(4)** 55.1 vs. **(5)** 23.1, $Z = 4.802$***	**(3)** 44.2 vs. **(4)** 54.1, $Z = 3.014$** **(3)** 44.2 vs. **(5)** 22.5, $Z = 4.986$*** **(4)** 54.1 vs. **(5)** 22.5, $Z = 5.262$***
Player 1 vs. Player 2 Contributions	**(3)** 73.4% vs. 72.2%, $Z = 0.380$ **(4)** 73.4% vs. 73.8%, $Z = 0.705$ **(5)** 62.9% vs. 39.7%, $Z = 4.234$***	**(3)** 73.8% vs. 73.1%, $Z = 0.414$ **(4)** 71.9% vs. 74.0%, $Z = 1.181$ **(5)** 61.2% vs. 39.0%, $Z = 4.575$***
Player 1 vs. Player 2 Payoffs	**(3)** 71.6 vs. 72.2, $Z = 0.366$ **(4)** 84.2 vs. 70.8, $Z = 3.239$** **(5)** 43.5 vs. 79.5, $Z = 5.464$***	**(3)** 71.9 vs. 72.3, $Z = 0.401$ **(4)** 84.3 vs. 69.8, $Z = 3.899$*** **(5)** 43.2 vs. 79.3, $Z = 5.736$***
First versus last **five rounds**	**(3)** 73.3% vs. 70.8%, $Z = 1.085$ **(4)** 73.2% vs. 74.6%, $Z = 0.079$ **(5)** 50.1% vs. 44.1%, $Z = 1.472$	**(3)** 73.7% vs. 72.5%, $Z = 1.034$ **(4)** 70.9% vs. 73.8%, $Z = 0.602$ **(5)** 50.3% vs. 42.8%, $Z = 2.049$*
Player 1 vs. Player 2 Reciprocity	**(3)** 83.6% vs. 84.3%, $Z = 0.492$ **(4)** 85.2% vs. 88.1%, $Z = 0.280$ **(5)** 85.9% vs. 44.8%, $Z = 4.252$***	**(3)** 83.6% vs. 84.2%, $Z = 0.454$ **(4)** 77.6% vs. 87.4%, $Z = 1.036$ **(5)** 86.8% vs. 43.1%, $Z = 4.648$***

Supplementary Table 2: Results for main and imputed samples under unequal productivities. Bold numbers in brackets refer to the respective treatment, **(3)** Unequal productivies, **(4)** Aligned inequality, **(5)** Misaligned inequality. The number of data points is $N_3 = 40$, $N_4 = 39$, $N_5 = 40$ (left column) and $N_3 = 43$, $N_4 = 44$ and $N_5 = 44$, respectively.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, using either a Mann-Whitney test (for between-treatment comparisons) or a Wilcoxon signed-rank test (for within-treatment comparisons). All tests are two-tailed.

According to the two tables, the main sample and the imputed sample yield qualitatively and quantitatively similar results. In fact, across all tests, there is only one difference. For the treatment with misaligned inequality, we have reported in the previous section that there is no significant trend in the players' contributions over time. Players contribute 50.1% of their endowment during the first five rounds of the experiment, whereas they contribute 44.1% during the last five rounds ($Z = 1.472$, $p = 0.141$). In contrast, for the imputed sample, the trend becomes significant. According to the imputed sample, contributions start at 50.3% during the first five rounds and drop to 42.8% during the last five rounds ($Z = 2.049$, $p = 0.041$, see **Supplementary Table 2**).

Overall this additional analysis based on imputed data suggests that the major conclusions of our experiment remain unchanged. When players have equal productivities, equal endowments result in the highest surplus. In contrast, when players differ in their productivities, the highest surplus is achieved under aligned inequality, where the most productive player receives a larger share of the endowment.

Robustness with respect to multiple testing. In the previous analyses, we have tested multiple hypotheses, and we have reported individual test statistics. Running multiple tests on the same data set increases the risk that some null hypotheses are rejected by chance, even if all hypotheses were true[78]. A commonly employed guard against such false discoveries is the Bonferroni correction, which makes an adjustment to the relevant threshold for the p-value, based on the number of hypotheses tested[79]. More specifically, if h is the number of tested hypotheses and α is the significance level of interest (e.g., $\alpha = 0.05$ or $\alpha = 0.01$), then the Bonferroni corrected significance level becomes α/h.

In the following, we explore which of our key results reported in the main text remain valid if we use the Bonferroni method to correct for multiple testing. In **Fig. 4**, we explore the impact of inequality on two quantities, total contributions and generated surplus. To explore whether there are significant differences between treatments for these two quantities, we have performed eight statistical tests in total (see the first two rows in **Supplementary Table 1** and **Supplementary Table 2**). As a consequence, the Bonferroni corrected significance levels become $0.05/8 = 0.00625$ (1 asterisk), $0.01/8 = 0.00125$ (2 asterisks), and $0.001/8 = 0.000125$ (3 asterisks).

Using these stricter significance levels, we find that all our main text conclusions remain valid after applying the Bonferroni correction. For example, for equal productivities we have found that unequal endowments lead to lower contributions and a reduced surplus. The corresponding p-value of $p = 0.005$ remains significant even if we correct for multiple testing (although the relevant significance level is now $0.05/8$, i.e., one asterisk instead of two). Similarly, we also find for the treatments with unequal productivities that all reported differences in the players' contributions and their surplus remain significant (although the number of asterisks may again change).

In sum, this analysis suggests that our key findings reported in the main text remain unchanged if we account for multiple testing.

6   Discussion

In this last section, we close by critically reviewing some of our model assumptions, and by suggesting a few model extensions for future work.

On the rigidity of endowments and productivities. In the main text, we have discussed our framework by focusing on a linear public good game where players may differ in their endowments and their productivities. When comparing the players' contributions across different public good games, we have often tacitly assumed that a player's productivities are given, whereas their endowments can be varied. There are two reasons why we consider this assumption to be useful.

First, we believe that in many applications, the players' endowments are best interpreted as their monetary resources, whereas their productivities could reflect their skills. In such a scenario, it is plausible to assume that in the short run, endowments are more amenable to interventions than productivities.

The second reason is of a more conceptual nature. When comparing the outcome of different public good games, one needs to make sure that games are commensurable. In the case of endowments, there exists a natural linear budget constraint: after any endowment redistribution, the sum of total endowments needs to remain unchanged. This budget constraint allows us to compare the consequences of different endowment distributions (provided they all sum up to the same value). With respect to the players' productivities, it is somewhat less clear what the appropriate budget constraint would be. In some cases it may again be natural to assume that the sum over all players' productivities is constant. In some other cases, however, it may be more natural to assume that it becomes disproportionally difficult to further increase the productivity of individuals who already are highly productive. To allow for a clean comparison, we have thus decided to take the players' productivities as fixed, and to vary their endowments. However, our framework and our analytical results are general. Once a budget constraint for the players' productivities is specified, our theoretical results immediately apply.

Public good contributions and taxation. In our model we do not consider the process by which endowments could be redistributed among players. Instead, we take different distributions as given and explore their impact on cooperation and welfare. However, there seem to be natural parallels between our framework and the question of how a social planner should set the tax rate in a population. This social planner might collect some of the group members' endowments and redistribute them. Alternatively, the social planner might directly invest the collected endowments into the public good.

It is not clear how this latter mechanism could be mapped onto our framework. As a first approximation, one could assume that in addition to the players $1,...,n$ there is another player 0 (the social planner). Player 0 has no own endowment. Instead Player 0 determines a tax rate T, which all other group members need to pay. Depending on the implementation of the tax, it may either be a player's total endowment that is subject to taxation, or only the proportion of the endowment the player decided to withhold. Given the implemented tax, all players $1,...,n$ would then decide how much to contribute to

the public good. Again, these contributions would be multiplied by some individual factor r_i. In addition, the total amount of tax collected is multiplied by r_0. The resulting sum is then equally distributed among all n group members (excluding Player 0).

In this extended framework, the problem of optimal taxation can be studied as a problem of mechanism design. How should the social planner set the tax rate T, and which implementation should be chosen, in order to maximize group payoffs? A particularly interesting case might arise if the social planner's own productivity r_0 is considerably below the maximum productivity $\max_i r_i$ of the group members.

However, we note that the above account of taxation abstracts away several important aspects. For example, it assumes that players and the social planner would routinely contribute to the same public good. In many applications it seems instead more natural that individuals can choose between different public goods (e.g., creating new jobs versus donating money to universities). This choice might in turn depend on which public good is most relevant to the respective individual. Due to these intricacies, we have refrained from modeling taxation explicitly, and leave it for future research instead. A related question is which kind of taxation a population considers as fair, which hinges heavily on perceptions of income inequality in a society[80].

Central institutions and corruption. Herein, we have explored how cooperation in asymmetric groups can emerge when individuals interact repeatedly. In our model, individuals cooperate because there is a shadow of the future. If subjects were to withhold their contributions, they risk that also other group members reduce their contributions in subsequent rounds.

Instead, there is also a rich theoretical and experimental literature that explores the effects of centralized or decentralized punishment institutions on cooperation[81–87]. A few of these experimental studies explicitly take into account that players may be heterogeneous[29,36,39]. However, it seems to us that a full theoretical understanding of how punishment institutions act in asymmetric groups is still missing.

In this context, it may also be interesting to study how the efficiency of central punishment institutions is undermined by bribery and corruption[88]. In the context of homogeneous groups, corruption possibilities can lead to a large decrease in public good provisioning, with the exact welfare losses depending on the cultural background of a society, the implemented transparency rules, and the extent to which leaders themselves are invested in the public good[89]. In the context of heterogeneous groups, corruption may have further detrimental implications. For example, when players with a higher endowment are better capable of bribing public sanctioning institutions, any efficiency advantages of unequal endowments may be dampened or reversed. While our present framework does not address these issues, we believe these questions are worthwhile of further study.

7 Appendix

7.1 Proofs

Proof of Proposition 1.

(1) ⇒ (2) Suppose $\mathbf{e} \in E_u(\delta)$. Then there is a strategy for each player such that no player has an incentive to deviate by not cooperating in every round, for the given endowment distribution $\mathbf{e}$. In the following, we refer to players who play their equilibrium strategy as 'residents', and to players who deviate as 'mutants'. If all players are residents, the equilibrium payoff is $\pi_R = u_i(\mathbf{e}, \mathbf{1})$.

Now suppose player i considers deviating by using the strategy ALLD instead. Then her payoff in the first round is $u_i(\mathbf{e}, \mathbf{1}_{-i})$. Due to Property **(PE)**, her payoff in all subsequent rounds is at least $u_i(\mathbf{e}, \mathbf{0})$. Hence, the mutant payoff to a deviating ALLD player is at least

$$\pi_M \geq (1-\delta)u_i(\mathbf{e}, \mathbf{1}_{-i}) + \delta u_i(\mathbf{e}, \mathbf{0}). \tag{31}$$

Since the strategies of the residents form a subgame perfect equilibrium, we can conclude that

$$(1-\delta)u_i(\mathbf{e}, \mathbf{1}_{-i}) + \delta u_i(\mathbf{e}, \mathbf{0}) \leq u_i(\mathbf{e}, \mathbf{1}), \tag{32}$$

or equivalently,

$$\delta\Big(u_i(\mathbf{e}, \mathbf{1}_{-i}) - u_i(\mathbf{e}, \mathbf{0})\Big) \geq u_i(\mathbf{e}, \mathbf{1}_{-i}) - u_i(\mathbf{e}, \mathbf{1}). \tag{33}$$

(2) ⇒ (3) To prove that Grim is a subgame perfect equilibrium, we can make use of the one-shot deviation principle[44]. That is, we only need to show that no mutant can get a higher payoff if she deviates in one round, and acts exactly as the other residents in all subsequent rounds. To this end, let us first assume that we are in the first round, or that all players have contributed their full endowment in all previous rounds. As the residents apply the strategy Grim, they continue to contribute their full endowment in the next round. A one-shot deviation consists in contributing an amount lower than the full endowment. In that case, Property **(IF)** implies that the maximum one-shot deviation payoff can be achieved by not contributing at all, in which case the continuation payoff becomes

$$\pi_M = (1-\delta)u_i(\mathbf{e}, \mathbf{1}_{-i}) + \delta u_i(\mathbf{e}, \mathbf{0}). \tag{34}$$

Given the assumptions in (2), this continuation payoff is at most the resident payoff $u_i(\mathbf{e}, \mathbf{1})$.

Alternatively, let us now assume that players did not contribute their full endowment in all past rounds. In that case, a one-shot deviation requires player i to contribute a positive amount in the next rounds, whereas all other players continue to contribute nothing. By Property **(IF)**, such a deviation is unprofitable.

(3) ⇒ (1) Since in a group of Grim players everyone contributes her full endowment in every round, and since Grim forms a subgame perfect equilibrium by assumption, full cooperation is feasible. □

Proof of Proposition 2.

1. Suppose $\mathbf{e} \in E_u(\delta)$. Because of Proposition 1, it follows that

$$\delta\big(u_i(\mathbf{e}, \mathbf{1}_{-i}) - u_i(\mathbf{e}, \mathbf{0})\big) \geq u_i(\mathbf{e}, \mathbf{1}_{-i}) - u_i(\mathbf{e}, \mathbf{1}).$$

   Since $\delta' > \delta$ by assumption and $u_i(\mathbf{e}, \mathbf{1}_{-i}) \geq u_i(\mathbf{e}, \mathbf{0})$ because of Property **(PE)**, it follows that

$$\delta'\big(u_i(\mathbf{e}, \mathbf{1}_{-i}) - u_i(\mathbf{e}, \mathbf{0})\big) \geq \delta\big(u_i(\mathbf{e}, \mathbf{1}_{-i}) - u_i(\mathbf{e}, \mathbf{0})\big) \geq u_i(\mathbf{e}, \mathbf{1}_{-i}) - u_i(\mathbf{e}, \mathbf{1}).$$

   Therefore, again by Proposition 1, $\mathbf{e} \in E_u(\delta')$.

2. Consider the map $\varphi : \Delta^n \to \Delta^n$, where Δ^n is the n-dimensional unit simplex, defined by

$$\varphi(\mathbf{e}) = \frac{u(\mathbf{e}, \mathbf{1})}{U(\mathbf{e}, \mathbf{1})}.$$

   Because $u(\mathbf{e}, \mathbf{x})$ is continuous in $\mathbf{e}$ due to property **(C)**, and because $U(\mathbf{e}, \mathbf{1}) > U(\mathbf{e}, \mathbf{0}) = 1$ due to property **(OC)**, we conclude that φ is continuous. Since Δ^n is compact and convex, it follows by Brouwer's fixed point theorem that there exists an endowment distribution $\mathbf{e}' = (e_1', \ldots, e_n')$ such that $\varphi(\mathbf{e}') = \mathbf{e}'$. As a consequence, we obtain that for all players i the following relationship holds,

$$u_i(\mathbf{e}', \mathbf{1}) \;=\; U(\mathbf{e}', \mathbf{1}) \cdot e_i' \;>\; e_i' \;=\; u_i(\mathbf{e}', \mathbf{0}). \tag{35}$$

   Let us now define δ' by

$$\delta' = \max_i \left\{ \frac{u_i(\mathbf{e}', \mathbf{1}_{-i}) - u_i(\mathbf{e}', \mathbf{1})}{u_i(\mathbf{e}', \mathbf{1}_{-i}) - u_i(\mathbf{e}', \mathbf{0})} \right\}. \tag{36}$$

   By inequality (35), $\delta' < 1$. Moreover, $\mathbf{e}' \in E_u(\delta')$ because the inequalities in condition (6) are satisfied by construction of δ'. The claim that $E_u(\delta) \neq \emptyset$ for all $\delta \geq \delta'$ then follows from the first part of this proof. $\qquad\square$

Proof of Proposition 3. For a public good game with payoff function u and continuation probability δ, we define the following function with respect to player i,

$$F_i(\mathbf{e}) = \delta\Big(u_i(\mathbf{e}, \mathbf{1}_{-i}) - u_i(\mathbf{e}, \mathbf{0})\Big) - \Big(u_i(\mathbf{e}, \mathbf{1}_{-i}) - u_i(\mathbf{e}, \mathbf{1})\Big). \tag{37}$$

According to Proposition 1, full cooperation is feasible for some given endowment distribution $\mathbf{e}$ if and only if $F_i(\mathbf{e}) \geq 0$ for all i. In the special case that player i is given the whole endowment (that is, $\mathbf{e}$ is such that $e_i = 1$ and $e_j = 0$ for all $j \neq i$), it follows from Property **(PE)** that $u_i(\mathbf{e}, \mathbf{1}_{-i}) = u_i(\mathbf{e}, \mathbf{0})$, whereas it

follows from Property **(IF)** that $u_i(\mathbf{e}, \mathbf{1}_{-i}) > u_i(\mathbf{e}, \mathbf{1})$. For that endowment distribution $\mathbf{e}$ we thus obtain $F_i(\mathbf{e}) < 0$. Since $u(\mathbf{e}, \mathbf{x})$ is continuous in $\mathbf{e}$ because of Property **(C)**, $F_i(\mathbf{e})$ is continuous as well. Hence there is an $\varepsilon_i > 0$ such that $F_i(\mathbf{e}') < 0$ for all $\mathbf{e}'$ for which $e_i' > 1 - \varepsilon_i$. That is, for all such endowment distributions $\mathbf{e}'$, cooperation is infeasible. □

The proofs of Propositions 4 and 5 is by example, provided in Sections 3.2 and 3.3. The proof of Proposition 6 is by iterated application of the following Lemma.

Lemma. *Consider a standard public goods game with payoffs according to Eq. (10), and assume productivities satisfy $r_1 < \ldots < r_n$. Let $\mathbf{e}^1$ be an endowment vector for which endowments are not ordered according to the players' productivities, such that there is an index k such that $e_k^1 > e_{k+1}^1$. Let $\mathbf{e}^2$ be the endowment vector for which $e_k^2 = e_{k+1}^1$, $e_{k+1}^2 = e_k^1$ and $e_j^2 = e_j^1$ for all other j. If full cooperation is feasible for $\mathbf{e}^1$ then so it is for $\mathbf{e}^2$.*

Proof. Rewriting the condition (11), it follows that cooperation is feasible for a given endowment distribution $\mathbf{e}$ if and only if

$$\big(n - (1-\delta)r_i\big)e_i \le \delta \sum_{j=1}^{n} r_j e_j \tag{38}$$

for all players i. Equivalently, if we define the two auxiliary functions $G(\mathbf{e}) := \max_i \big\{\big(n - (1-\delta)r_i\big)e_i\big\}$ and $H(\mathbf{e}) := \delta \sum_{j=1}^{n} r_j e_j$, cooperation is feasible if and only if $G(\mathbf{e}) \le H(\mathbf{e})$. Now, using the properties of the two given endowment vectors $\mathbf{e}^1$ and $\mathbf{e}^2$, we have

$$\begin{aligned}
\big(n - (1-\delta)r_k\big)e_k^2 &= \big(n - (1-\delta)r_k\big)e_{k+1}^1 < \big(n - (1-\delta)r_k\big)e_k^1 \\
\big(n - (1-\delta)r_{k+1}\big)e_{k+1}^2 &= \big(n - (1-\delta)r_{k+1}\big)e_k^1 \le \big(n - (1-\delta)r_k\big)e_k^1,
\end{aligned}$$

and

$$\big(n - (1-\delta)r_j\big)e_j^2 = \big(n - (1-\delta)r_j\big)e_j^1$$

for all other $j \notin \{k, k+1\}$. Thus, we can conclude

$$G(\mathbf{e}^2) \le G(\mathbf{e}^1). \tag{39}$$

On the other hand, we have

$$H(\mathbf{e}^1) = \delta \sum_{j \notin \{k,k+1\}} r_j e_j^1 + \delta(r_k e_k^1 + r_{k+1} e_{k+1}^1) \le \delta \sum_{j \notin \{k,k+1\}} r_j e_j^1 + \delta(r_k e_{k+1}^1 + r_{k+1} e_k^1) = H(\mathbf{e}^2). \tag{40}$$

Thus, if cooperation is feasible for $\mathbf{e}^1$, and hence $G(\mathbf{e}^1) \le H(\mathbf{e}^1)$, it follows that

$$G(\mathbf{e}^2) \le G(\mathbf{e}^1) \le H(\mathbf{e}^1) \le H(\mathbf{e}^2). \tag{41}$$

That is, cooperation is feasible for the endowment vector $\mathbf{e}^2$. □

The proof of Proposition 7 is again by example, which is provided in Section 3.4. The proof of Proposition 8 is as follows.

Proof of Proposition 8. For arbitrary continuation probability δ and endowment distribution $\mathbf{e} = (e_1, e_2)$ it follows from Eq. (6) that $\mathbf{e} \in E_u(\delta)$ if and only if

$$
\begin{aligned}
\delta\, c_{12}\, e_2 &\geq -c_{11}\, e_1 \\
\delta\, c_{21}\, e_1 &\geq -c_{22}\, e_2.
\end{aligned}
\tag{42}
$$

Dividing both inequalities by e_2 (which is positive due to Proposition 3), we obtain that $\mathbf{e} \in E_u(\delta)$ if and only if

$$
\frac{-c_{22}}{\delta \cdot c_{21}} \leq \frac{e_1}{e_2} \leq \frac{\delta \cdot c_{12}}{-c_{11}}.
\tag{43}
$$

In particular, an endowment distribution $\mathbf{e} \in E_u(\delta)$ exists if and only if $-c_{22}/(\delta c_{21}) \leq \delta c_{12}/(-c_{11})$. This condition yields $\delta \geq \delta_u^*$, with δ_u^* as defined by the first equation in (21). Plugging this δ_u^* into either the left hand side or the right hand side of condition (43) shows that e_1/e_2 is given by the second equation in (21). □

Proof of Proposition 9. The proof is similar to the proof of Proposition 1, by employing the one-shot deviation principle[44]. That is, we show that no mutant can get a higher payoff by deviating for one round, and by reverting to WSLS thereafter. We consider two cases, depending on whether the history is such that the $n-1$ residents would cooperate (contribute their full endowment) or defect (contribute nothing). Let us first consider the case where players are either in the first round, or they have all chosen the same relative contribution in the previous round. As the residents use WSLS, they cooperate. By sticking to the resident strategy the mutant i obtains the usual resident payoff

$$
\pi_R = u_i(\mathbf{e}, \mathbf{1}).
\tag{44}
$$

If the mutant decides to deviate for one round instead, the best possible deviation is to contribute nothing at all. In that case the mutant's continuation payoff is

$$
\pi_M = (1-\delta)u_i(\mathbf{e}, \mathbf{1}_{-i}) + \delta\Big((1-\delta)u_i(\mathbf{e}, \mathbf{0}) + \delta u_i(\mathbf{e}, \mathbf{1})\Big).
\tag{45}
$$

For WSLS to be a subgame perfect equilibrium we require $\pi_R \geq \pi_M$, which yields the inequality (28).

As the second case, assume now that players differed in their last round's contributions. In that case, by adhering to the resident strategy WSLS player i obtains the continuation payoff

$$\pi_R = (1-\delta)u_i(\mathbf{e},\mathbf{0}) + \delta u_i(\mathbf{e},\mathbf{1}). \tag{46}$$

If the mutant instead contributes some positive amount in the next round, it follows from Property **(IF)** that her continuation payoff is at most

$$\pi_M = (1-\delta^2)u_i(\mathbf{e},\mathbf{0}) + \delta^2 u_i(\mathbf{e},\mathbf{1}). \tag{47}$$

We conclude that $\pi_R \geq \pi_M$ requires $u_i(\mathbf{e},\mathbf{1}) \geq u_i(\mathbf{e},\mathbf{0})$, which is a strictly weaker condition than inequality (28). $\qquad\square$

Proof of Proposition 10.

1. By Proposition 9, an endowment distribution $\mathbf{e}$ satisfies $\mathbf{e} \in E_u^W(\delta)$ if and only if

$$\delta \geq \frac{u_i(\mathbf{e},\mathbf{1}_{-i})-u_i(\mathbf{e},\mathbf{1})}{u_i(\mathbf{e},\mathbf{1})-u_i(\mathbf{e},\mathbf{0})} > 0 \quad \text{for all players } i \text{ with } e_i > 0. \tag{48}$$

   Because $\delta' > \delta$, $\mathbf{e} \in E_u^W(\delta)$ thus implies $\mathbf{e} \in E_u^W(\delta')$.

2. For a linear and symmetric public goods game with payoff function (7), Proposition 9 implies that $E_u^W(\delta) \neq \emptyset$ if and only if there is some endowment vector $\mathbf{e}$ such that

$$\max e_i \leq \frac{\delta c_1}{\delta c_1 + (1+\delta)c_2} \tag{49}$$

   Because $1/n \leq \max(e_i)$, it follows that $(1/n, \ldots, 1/n) \in E_u^W(\delta)$.

3. As a simple counter-example, consider the asymmetric payoff function (10) where players have different productivities. Proposition 9 implies that $\mathbf{e} \in E_u^W(\delta)$ if and only if

$$\delta\left(\frac{1}{n}\sum_{j=1}^n r_j e_j - e_i\right) \geq \left(1 - \frac{r_i}{n}\right) e_i \quad \text{for all players } i \text{ with } e_i > 0. \tag{50}$$

   For the 3-player case with $r_1 = 1.1$, $r_2 = 1.5$, $r_3 = 2.9$ and $\delta = 0.5$, a straightforward computation shows that condition (50) is violated for $\mathbf{e} = (1/3, 1/3, 1/3)$, but satisfied for $\mathbf{e} = (0, 1/5, 4/5)$.

4. For a fixed δ and $\mathbf{e} = (e_1, e_2)$, Proposition 9 implies that $\mathbf{e} \in E_u^W(\delta)$ if and only if the following two conditions hold,

$$\begin{aligned}
\delta\big(c_{11}e_1 + c_{12}e_2\big) &\geq -c_{11}e_1 \\
\delta\big(c_{21}e_1 + c_{22}e_2\big) &\geq -c_{22}e_2.
\end{aligned} \tag{51}$$

Dividing both inequalities by $e_2 > 0$ shows that $\mathbf{e} \in E_u^W(\delta)$ if and only if

$$\frac{-(1+\delta)c_{22}}{\delta c_{21}} \leq \frac{e_1}{e_2} \leq \frac{\delta c_{12}}{-(1+\delta)c_{11}}. \tag{52}$$

This chain of inequalities can only be met if the left hand's side is less than the right hand's side, which yields the following implicit condition for the minimal continuation probability for which $E_u^W(\delta)$ is non-empty,

$$\frac{\delta}{1+\delta} \geq \sqrt{\frac{c_{11}c_{22}}{c_{21}c_{12}}}. \tag{53}$$

By using this lower bound for $\delta/(1+\delta)$ and plugging it into (52), we obtain

$$\frac{e_1}{e_2} = \sqrt{\frac{c_{12}c_{22}}{c_{11}c_{21}}} \tag{54}$$

That is, we have shown that for WSLS to be an equilibrium for the lowest possible continuation probability, endowments need to meet condition (54), which coincides with the condition for $\mathbf{e}$ to be maximally cooperative. The statement then follows from the first part of this Proposition. $\square$

Proof of Proposition 11. The proof is analogous to the proof of Proposition 9, with the only difference that now all one-shot deviations are required to yield a strictly lower payoff. As a consequence, the inequality (29) needs to be strict, and all players are required to have a positive endowment, in order to exclude neutral one-shot deviations. The result then follows from the argument given by Boyd[62]. $\square$

7.2 MATLAB code used for the evolutionary computations

In the following, we provide the MATLAB algorithms that we have used to (*i*) calculate the payoffs of memory-1 players, (*ii*) simulate the introspection dynamics, and (*iii*) numerically compute the expected trajectories for introspection dynamics when players use pure strategies with two possible contributions, $X = \{0, 1\}$. (*iv*) Moreover, we provide our code for games where players use finite state automata to represent their strategies. For simplicity, the following code considers groups of size two, which has been used to create **Figs. 3–4** and **Extended Data Figs. 1–6** . The respective code for n-player games is similar, and is available from the authors upon request.

Algorithm to calculate the payoffs in a game between memory-1 players

```
function [pi,coop,v,COS,M]=payoff(P1,P2,e1,r1,r2,m);
% Output:
% pi, coop, v ..  payoffs, cooperation rates, and invariant distribution
% COS ..  List of possible outcomes of each round
% M ..  transition matrix of the Markov chain
% Input:
% P1 ..  m^2xm matrix that represents player 1's strategy;
%        First mxm matrix gives probability to choose contribution x1, etc.
% P2 ..  m^2xm matrix for player 2's strategy
% e1,r1,r2 ..  players' endowments and their productivities
% m ..  number of possible contribution levels, X={0,1/(m-1),...,1}

%% PARAMETERS AND PREPARATIONS %%
e2=1-e1;
pi=zeros(1,2); coop=zeros(1,2); % Initializing the output variables
OS=zeros(m^2,2); COS=zeros(m^2,2); PayOS=zeros(m^2,2);
% Possible outcomes/contributions/payoffs of the one-shot game
for i1=1:m % Going through all possible contribution levels for player 1
  for i2=1:m % Same for player 2
    j=(i1-1)*m+i2; % Storing i1, i2 in one variable that runs from 1 to m^2
    OS(j,:)=[i1,i2]; COS(j,:)=[(i1-1)/(m-1), (i2-1)/(m-1)];
    PayOS(j,:)=(e1*r1*COS(j,1)+e2*r2*COS(j,2))/2+[(1-COS(j,1))*e1,(1-COS(j,2))*e2];
  end
end

%% CONSTRUCTING THE MARKOV CHAIN %%
M=zeros(m^2,m^2); % Initializing the transition matrix
for i1=1:m, for i2=1:m % Running over the players' previous actions
  for j1=1:m, for j2=1:m % Running over players' next actions
    M((i1-1)*m+i2,(j1-1)*m+j2)=P1(i1+(j1-1)*m,i2) * P2(i1+(j2-1)*m,i2);
```

```
  end, end
end, end
v=null(M'-eye(m^2)); v=v/sum(v); % Calculating the normalized left eigenvector
pi=PayOS'*v; coop=COS'*v;
% Calculating cooperation rates and payoffs in the case of no discounting
end
```

Algorithm to simulate the introspection dynamics

```
function [AvCoop,AvPi,COS,Data]=SimulateIntrospection(e1,r1,r2,epsi,m,s,nGen);
% Output:
% AvCoop, AvPi ..  average cooperation and payoff over course of simulation
% COS ..  possible contributions in one-shot game
% Data ..  Summary of all parameters used
% Input:
% e1,r1,r2 ..  players' endowments and their productivities
% epsi ..  error rate
% m ..  number of possible contribution levels, X={0,1/(m-1),...,1}
% s ..  strength of selection
% nGen ..  number of generations (number of elementary updating events)

%% PARAMETERS AND PREPARATIONS %%
C=clock; rng(C(5)*60+C(6)); % Setting up the random number generator
P1=zeros(m^2,m); P2=zeros(m^2,m); % Initially, both players use ALLD
[pi,coop]=payoff(P1,P2,e1,r1,r2,m); % Calculate initial payoffs and cooperation
AvPi=pi; AvCoop=coop; % Initialize variables for respective averages
Data=['e1=',num2str(e1), '; r1=',num2str(r1), '; r2=',num2str(r2),...
  '; eps=',num2str(epsi), '; m=',num2str(m), '; s=',num2str(s),...
  '; nGen=',num2str(nGen)];

%% SIMULATING THE EVOLUTIONARY PROCESS %%
for t=2:nGen
 iPop=randi(2); % Randomly choosing a player to update
 Mut=getrandomstrategy(m,epsi); % Creating a mutant strategy for that player

 % Case 1:  Player 1 explores a new strategy
 if iPop==1
   [piM,coopM,vM,COS]=payoff(Mut,P2,e1,r1,r2,m); % Mutant strategy payoff
   rho=1/(1+exp(-s*(piM(1)-pi(1)))); % Player's updating probability
   if rand(1)<rho % If updating occurs
     P1=Mut; pi=piM; coop=coopM; % Update respective variables
```

```matlab
    end

  % Case 2:  Player 2 explores a new strategy
  else
    [piM,coopM,vM,COS]=payoff(P1,Mut,e1,r1,r2,m); % Same steps for player 2
    rho=1/(1+exp(-s*(piM(2)-pi(2))));
    if rand(1)<rho
      P2=Mut; pi=piM; coop=coopM;
    end
  end

  % Updating the averages
  AvPi=(t-1)/t*AvPi+1/t*pi;
  AvCoop=(t-1)/t*AvCoop+1/t*coop;
end
end


function Mut=getrandomstrategy(m,epsi);
% Auxiliary function that creates a sensible mutant strategy
Mut=zeros(m^2,m); % Initialize strategy
l=(0:m-1)*m; % Auxiliary vector
for i1=1:m % Running through all previous actions of player 1
  for i2=1:m % Running through previous actions of player 2
    v=rand(1,m-1); v=[0, v, 1]; v=sort(v); v=diff(v); Mut(i1+l,i2)=v;
    % Produces a random stochastic strategy by uniformly drawing a vector v
    % in the m-dimensional unit simplex
    % To produce a pure strategy, replace above line by following line
    % v=zeros(1,m); v(randi(m))=1; Mut(i1+l,i2)=v;
  end
end
Mut=(1-epsi)*Mut+epsi/(m-1)*(1-Mut); % Adding errors to the strategy
end
```

Algorithm to compute expected trajectories of introspection dynamics

```matlab
function [Coop1,Coop2,XStr1,XStr2,PayMat1,PayMat2,Str,Data]=...
    ExpectedIntrospectionDynamics(e1,r1,r2,epsi,s,nGen);
% Output:
% Coop1, Coop2 ..  The players' cooperation rates over time
```

```matlab
% XStr1, XStr2 ..  The players' strategies over time
% PayMat1, PayMat2 ..  Payoff matrices for each pair of memory-1 strategies
% Str ..  List of all memory-1 strategies in notation (pCC,pCD,pDC,pDD)
% Data ..  List of the parameters used
% Input:
% e1,r1,r2 ..  players' endowments and their productivities
% epsi ..  error rate
% s ..  strength of selection
% nGen ..  number of generations (number of elementary updating events)

%% PARAMETERS AND PREPARATIONS %%
e2=1-e1;
POS1=[(r1*e1+r2*e2)/2,r1*e1/2,r2*e2/2+e1,e1]; % One-shot payoffs for player 1
POS2=[(r1*e1+r2*e2)/2,r1*e1/2+e2,r2*e2/2,e2]; % One-shot payoffs for player 2
Str=[0 0 0 0; 0 0 0 1; 0 0 1 0; 0 0 1 1; 0 1 0 0; 0 1 0 1; 0 1 1 0; 0 1 1 1;
 1 0 0 0; 1 0 0 1; 1 0 1 0; 1 0 1 1; 1 1 0 0; 1 1 0 1; 1 1 1 0; 1 1 1 1];
% List of all 16 memory-1 strategies with contributions either x=0 or x=1.
ns=size(Str,1); % number of strategies
Data=['e1=',num2str(e1), '; r1=',num2str(r1), '; r2=',num2str(r2),...
 '; eps=',num2str(epsi), '; s=',num2str(s), '; nGen=',num2str(nGen)];

%% CREATING PAIRWISE PAYOFF MATRICES FOR EACH PAIR OF MEMORY-1 STRATEGIES %%
PayMat1=zeros(ns,ns); PayMat2=zeros(ns,ns); % Initializing the payoff matrices
CopMat1=zeros(ns,ns); CopMat2=zeros(ns,ns); % Matrices for cooperation rates
for i1=1:ns % Running through all strategies of player 1
  for i2=1:ns % Running through all strategies of player 2
    s1=Str(i1,:); s2=Str(i2,:);
    [pi1,pi2,coop1,coop2]=CalcPay(s1,s2,POS1,POS2,epsi);
    PayMat1(i1,i2)=pi1; PayMat2(i1,i2)=pi2;
    CopMat1(i1,i2)=coop1; CopMat2(i1,i2)=coop2;
  end
end

%% CREATING THE TRANSITION MATRIX W %%
W=zeros(ns^2, ns^2); % State space:  all ns^2 combinations of strategies
for i1Old=1:ns % Previous strategy of player 1
  for i2Old=1:ns % Previous strategy of player 2
    pi1Old=PayMat1(i1Old,i2Old); pi2Old=PayMat2(i1Old,i2Old); % Previous payoffs

    % Case 1:  Transitions where player 1 is chosen to update
    for i1New=1:ns % Running through possible new strategies for player 1
      if i1New~=i1Old
        pi1New=PayMat1(i1New,i2Old); % New payoff
```

45

```matlab
      rho=1/(1+exp(-s*(pi1New-pi1Old))); % Update probability
      W(ns*(i1Old-1)+i2Old,ns*(i1New-1)+i2Old)=1/(2*(ns-1))*rho;
    end
  end


  % Case 2:  Transitions where player 2 is chosen to update
  for i2New=1:ns % Running through possible new strategies for player 2
   if i2New~=i2Old
     pi2New=PayMat2(i1Old,i2New);
     rho=1/(1+exp(-s*(pi2New-pi2Old)));
     W(ns*(i1Old-1)+i2Old,ns*(i1Old-1)+i2New)=1/(2*(ns-1))*rho;
   end
  end
 end
end
for k=1:ns^2
 W(k,k)=1-sum(W(k,:)); % Defining diagonal entries such that total sum is 1
end

%% CALCULATING TRAJECTORIES OVER TIME %%
Coop1=zeros(1,nGen); Coop2=zeros(1,nGen); % Initializing the cooperation vectors
XStr1=zeros(ns,nGen); XStr2=zeros(ns,nGen); % List of the strategies over time
coop1v=zeros(ns^2,1); coop2v=zeros(ns^2,1);
i1s=zeros(ns,ns); i2s=zeros(ns,ns); % Matrices that indicate which entries
 % of W correspond to player 1's strategy i1s (row) and 2's strategy i2s (column)
for i=1:ns
 i1s(i,:)=ns*(i-1)+1:ns*i;
 i2s(i,:)=(0:ns-1)*ns+i;
end
for i=1:ns % Transforming the possible cooperation rates into a vector
 coop1v(ns*(i-1)+1:ns*i,:)=CopMat1(i,:);
 coop2v(ns*(i-1)+1:ns*i,:)=CopMat2(i,:);
end
v=zeros(1,ns^2); v(1)=1; % Initially everyone plays ALLD
for t=1:nGen % Calculating the expected trajectory for each time step
 v=v*W; % Iterate process for one time step
 Coop1(t)=v*coop1v; Coop2(t)=v*coop2v; % Expected cooperation rate at time t
 for i=1:ns
  XStr1(i,t)=sum(v(i1s(i,:))); XStr2(i,t)=sum(v(i2s(i,:)));
  % Expected frequency of each strategy at time t
 end
end
end
```

```matlab
function [pi1,pi2,coop1,coop2]=CalcPay(p,q,POS1,POS2,epsi);
% Subroutine that calculates payoffs and cooperation rates for two players
% with memory-1 strategies p=(pCC,pCD,pDC,pDD) and q=(qCC,qCD,qDC,qDD)
p=p*(1-epsi)+(1-p)*epsi; q=q*(1-epsi)+(1-q)*epsi; % Adding noise
% Creating the transition matrix M
M=[p(1)*q(1), p(1)*(1-q(1)), (1-p(1))*q(1), (1-p(1))*(1-q(1));
   p(2)*q(3), p(2)*(1-q(3)), (1-p(2))*q(3), (1-p(2))*(1-q(3));
   p(3)*q(2), p(3)*(1-q(2)), (1-p(3))*q(2), (1-p(3))*(1-q(2));
   p(4)*q(4), p(4)*(1-q(4)), (1-p(4))*q(4), (1-p(4))*(1-q(4))];
v=null(M'-eye(4)); v=v/sum(v); % v ..  normalized left eigenvector of M
pi1=POS1*v; pi2=POS2*v; coop1=v(1)+v(2); coop2=v(1)+v(3);
end
```

Algorithm when strategies are represented by finite state automata

```matlab
function [AvgC,AvgPay,Data]=EvolProcFSA(rvec,evec,nCL,s,nGen);
% Output:
% AvgC, AvgPay ..  average cooperation and payoff over course of simulation
% Data ..  string that stores the used parameter values
% Input:
% rvec=(r1,r2), evec=(e1,e2) ..  productivity and endowment vectors
% nCL ..  number of different contribution levels (e.g.  2, 3, or 4).
% s ..  strength of selection, nGen ..  number of updating events.

%% PARAMETERS AND PREPARATIONS
S1=ones(1,nCL+1); S2=S1; % Initially, both players use ALLD (choose first contribution
level, always go to state 1)
n1=1; n2=1; % current number of states of the two finite state automata (FSA)
[pi,coop]=CalcPayoff(S1,S2,rvec,evec,nCL);
AvgPay=pi; AvgC=coop;
Data=['e1=',num2str(evec(1)), '; r1=',num2str(rvec(1)), '; r2=',num2str(rvec(2)),...
  '; nCL=',num2str(nCL), '; s=',num2str(s), '; nGen=',num2str(nGen)];

%% EVOLUTIONARY PROCESS
for i=2:nGen
 % Randomly choosing a player
 iPop=randi(2);

 % Player 1 explores a new strategy
```

```matlab
  if iPop==1
    Mut=getrandomstrategy(S1,nCL);
    [piM,coopM]=CalcPayoff(Mut,S2,rvec,evec,nCL);
    rho=1/(1+exp(-s*(piM(1)-pi(1))));
    if rand(1)<rho
      S1=Mut; pi=piM; coop=coopM;
    end

  % Player 2 explores a new strategy
  elseif iPop==2
    Mut=getrandomstrategy(S2,nCL);
    [piM,coopM]=CalcPayoff(S1,Mut,rvec,evec,nCL);
    rho=1/(1+exp(-s*(piM(2)-pi(2))));
    if rand(1)<rho
      S2=Mut; pi=piM; coop=coopM;
    end
  end

  % Updating the averages
  AvgPay=(i-1)/i*AvgPay+1/i*pi;
  AvgC=(i-1)/i*AvgC+1/i*coop;
end
end

function Mut=getrandomstrategy(S,nCL)
% Auxiliary program that creates a random mutant based on a given strategy S
% Strategies take the form of an ls x (nCL+1) matrix
% ls is the number of states of the FSA
% The first column defines which action to take in respective state
% The other nCL entries in each row specify the next state of the FSA,
% depending on which of the nCL possible actions the co-player has chosen

%% SETTING THE PARAMETERS
qAdd=0.3; % Probability to add a state
qRemove=0.35; % Probability to remove a state
qAction=0.175; % Probability to change an action
qArrow=0.175; % Probability to change an arrow
lS=size(S,1); % Number of states of the original FSA
Mut=S;

%% CREATING THE MUTANT
qRand=rand(1);
```

```matlab
% Adding a state
if qRand<qAdd
 Mut(lS+1,:)=[randi(nCL),randi(lS+1,1,nCL)]; % Append new state to end of Mut
 Mut(randi(lS),1+randi(nCL))=lS+1; % Choose a random arrow to point to new state
end

% Deleting a state
if qRand>=qAdd & qRand<qAdd+qRemove & lS>1
 iD=randi(lS); % Choosing a random state to be removed
 Act=Mut(:,1); Trans=Mut(:,2:end); % Distinguishing between actions and transitions
 [j1,j2]=find(Trans==iD); % Find all arrows that point to the state to be removed
 for i=1:length(j1)
   Trans(j1(i),j2(i))=-1; % Overwriting all arrows that point to removed state
 end
 for i=iD+1:lS
   [j1,j2]=find(Trans==i); % Find arrows that point to state with higher index
   for l=1:length(j1)
     Trans(j1(l),j2(l))=i-1; % Decrease index of those arrows by one
   end
 end
 [j1,j2]=find(Trans==-1);
 for i=1:length(j1) % Assign new target states for arrows to removed state
   Trans(j1(i),j2(i))=randi(lS-1);
 end
 Mut=[Act'; Trans']'; Mut(iD,:)=[]; % Removing the state with index iD
end

% Changing an action
if qRand>=qAdd+qRemove & qRand<qAdd+qRemove+qAction
 rands=randi(lS); % Choosing a random state to change its action
 as=S(rands,1);
 while as==S(rands,1);
   as=randi(nCL); % Assigning a new action to respective state
 end
 Mut(rands)=as;
end

% Changing an arrow
if qRand >= qAdd+qRemove+qAction & lS>1
 i=randi(lS); j=1+randi(nCL); % Choosing a random arrow to be changed
 ar=Mut(i,j);
 while ar==Mut(i,j)
   Mut(i,j)=randi(lS); % Assigning new target state to respective arrow
```

```matlab
    end
  end
end

function [pi,coop]=CalcPayoff(S1,S2,rvec,evec,nCL);
% Auxiliary program that computes the payoffs and cooperation rates
% when two players with FSA strategies S1 and S2 interact.
% rvec=[r1,r2] and evec=[e1,e2] contain the players' productivities and endowments
% nCL is the number of different contribution levels

%% PREPARATIONS
l1=size(S1,1); l2=size(S2,1); % number of states of the FSA of the two players
lContr=zeros(1,nCL); for i=1:nCL, lContr(i)=(i-1)/(nCL-1); end
% List of possible contributions (in percentage of the player's endowment
svec=zeros(l1*l2,2); cvec=svec; pivec=svec;
% Lists that store the players' visited states, their contributions,
% and their payoffs in the first l1*l2 rounds

%% COMPUTING THE OUTPUT
stop=0; j=0; % Initializing a stop variable and a counter
cS1=1; cS2=1; % Initial states of the two players
while stop==0
  j=j+1;   a1=S1(cS1,1); a2=S2(cS2,1); % Player's actions in the respective round
  svec(j,:)=[cS1,cS2]; % Storing the player's current states
  cvec(j,:)=[lContr(a1),lContr(a2)]; % Storing their current contributions
  pivec(j,:)=(rvec(1)*lContr(a1)*evec(1)+rvec(2)*lContr(a2)*evec(2))/2+...
    +[(1-lContr(a1))*evec(1),(1-lContr(a2))*evec(2)];
  % Storing the players' current payoffs
  % Check whether current state combination has already occurred earlier:
  for k=1:j-1
    if svec(k,:)==svec(j,:), stop=1; loopstart=k; end
  end
  cS1=S1(cS1,1+a2); cS2=S2(cS2,1+a1);
  % Compute the next states of the FSA of the two players
end
coop=mean(cvec(loopstart+1:j,:),1); pi=mean(pivec(loopstart+1:j,:),1);
% Calculate average cooperation rates and payoffs by averaging over all
% values in the loop that the two players will reach eventually
end
```

7.3 Game instructions of the behavioral experiment

In the following, we provide screenshots of the 'aligned inequality treatment' from the perspective of player 1. The instructions for player 2 and for the other treatments are analogous. They are available from the authors upon request.

SUMMARY

- You will be matched with another worker who is also taking this HIT at the same time.

- You will interact in real-time with this worker.

- There will be many interactive rounds with the same worker.

- The decisions you and the other worker make will affect your bonus payment.

Depending on the decisions you and the other worker make, you can earn a bonus between $0.00 and $2.59. However, you must **stay on task at all times**, otherwise you will automatically be disconnected and earn $0.00.

Detailed Instructions (1/3)

In this HIT, you are matched with another worker who is also taking this HIT. You will interact with the same worker for the entire task. One of you will be **Player 1** and the other will be **Player 2**. All interactions take place in real time.

Please be considerate of the worker's time!

In this task, there are multiple rounds. At the beginning of every round, each player receives a certain number of units (income). The total number of units distributed in each round will be 100 units: Player 1 will receive 75 units in every round, while Player 2 will receive 25 units in every round.

There will be at least 20 rounds of the group task. After the 20th round, there will be a 50% chance that another round will follow. After that round, there will again be a 50% chance that another round will follow, and so on until the group task ends. When the group task is over, you will see a brief questionnaire before you can submit the HIT.

Please answer the following questions:

What are the incomes of Players 1 and 2 in every round?

○ **Players 1 and 2 each receive 50 units every round.**

○ **Player 1 receives 75 units per round and Player 2 receives 25 units per round.**

○ **Players 1 and 2 receive different units each round.**

What is the chance that there will be another round **after the 20th round**?

○ **0%**

○ **50%**

○ **100%**

What is the chance that there will be another round **after the 21th round**?

○ **0%**

○ **50%**

○ **100%**

Detailed Instructions (2/3)

In every round, you have to decide how much of your income (in units) you want to contribute to the common group project and how much to keep for yourself.

All units that are contributed to the common group project are multiplied: every unit that **Player 1** contributed is multiplied by a **factor of 1.9** while every unit that **Player 2** contributed is multiplied by a **factor of 1.3**.

Once multiplied, the total amount is **split evenly** among both players, regardless of whether they contributed. Thus for every unit players contribute, they personally get less than a unit back: so no matter what the other player contributes, **you personally lose units on contributing, but contributing benefits the group as a whole.**

The exchange rate is 800 units = $1.00.

Please answer the following questions:

By what value will **Player 1's** contribution be multiplied?

○ **1.3**

○ **1.6**

○ **1.9**

By what value will **Player 2's** contribution be multiplied?

○ **1.3**

○ **1.6**

○ **1.9**

What level of contribution earns the **most money for the group as a whole?**

○ **None of your income.**

○ **Half of your income.**

○ **All of your income.**

What level of contribution earns the **most money for you personally?**

○ **None of your income.**

○ **Half of your income.**

○ **All of your income.**

Detailed Instructions (3/3)

These two **examples** illustrate how payoffs are calculated every round. Example A demonstrates that the highest payoff for the group is achieved if both contributed all their units to the common group project. Example B shows that a player will always earn more personally when contributing nothing.

EXAMPLE A:

	Individual players			Common group project	
	Player 1	Player 2		Due to Player 1	Due to Player 2
Initial endowment	75 units	25 units			
	-75 units	-25 units	→ Contributions	75 units	25 units
			Multiplier	× 1.9	× 1.3
			Multiplied contributions	142.5 units	32.5 units
Income from group project	87.5 units	87.5 units	← Total value (split evenly)	175 units	
TOTAL EARNINGS	**87.5 units**	**87.5 units**			

EXAMPLE B:

	Individual players			Common group project	
	Player 1	Player 2		Due to Player 1	Due to Player 2
Initial endowment	75 units	25 units			
	-65 units	0 units	→ Contributions	65 units	0 units
			Multiplier	× 1.9	× 1.3
			Multiplied contributions	123.5 units	0 units
Income from group project	61.8 units	61.8 units	← Total value (split evenly)	123.5 units	
TOTAL EARNINGS	**71.8 units**	**86.8 units**			

Please answer the following questions:

In **Example A**, what is **Player 1's income from the common group project**?

○ **25.5 units**

○ **70 units**

○ **87.5 units**

	Player	**Income**	**Multiplier**
YOU >>>	Player 1	75 units	1.9x
	Player 2	25 units	1.3x

Your income is 75 units.

Please decide now how many units you want to contribute to the common pool.

[   ]

(Type a value between 0 and 75.)

[Submit ...]

54

Player	Income	Multiplier	Contribution	Multiplied contribution
YOU >>> Player 1	75 units	1.9x	50 units	**95.00 units**
Player 2	25 units	1.3x	20 units	**26.00 units**
				Total: 121.00 units
				Share per person: 60.50 units

YOUR PAYOFF IN THIS ROUND:

Units you did not contribute:	**25.00** units
Units you earned from common pool:	**60.50** units
TOTAL THIS ROUND:	**85.50 units**

Continue ...

Thank you for playing!

You just finished playing the last round of the game.

In every game, there are at least 20 rounds. After that, another round is played with a 50% probability. And after that, yet another round is plyed with a 50% probability and so on. The random generator took this into account and determined a maximum length of 21 rounds for your group. You just finished with the 21th round and you can now claim your bonus.

Please finish the survey to receive your completion fee and your bonus. Your bonus is the number of units that you accumulated over all rounds.

Continue ...

8 Supplementary References

[1] Kerr, B., Godfrey-Smith, P. & Feldman, M. W. What is altruism? *Trends in Ecology & Evolution* **19**, 135–140 (2004).

[2] Nowak, M. A. Evolving cooperation. *Journal of Theoretical Biology* **299**, 1–8 (2012).

[3] Milinski, M., Sommerfeld, R. D., Krambeck, H.-J., Reed, F. A. & Marotzke, J. The collective-risk social dilemma and the prevention of simulated dangerous climate change. *Proceedings of the National Academy of Sciences USA* **105**, 2291–2294 (2008).

[4] Pacheco, J. M., Santos, F. C., Souza, M. O. & Skyrms, B. Evolutionary dynamics of collective action in n-person stag hunt dilemmas. *Proceedings of the Royal Society B* **276**, 315–321 (2009).

[5] Archetti, M. & Scheuring, I. Review: Game theory of public goods in one-shot social dilemmas without assortment. *Journal of Theoretical Biology* **299**, 9–20 (2012).

[6] Abou Chakra, M. & Traulsen, A. Evolutionary dynamics of strategic behavior in a collective-risk dilemma. *PLoS Computational Biology* **8**, e1002652 (2012).

[7] Jacquet, J. *et al.* Intra- and intergenerational discounting in the climate game. *Nature Climate Change* **3**, 1025–1028 (2013).

[8] Abou Chakra, M., Bumann, S., Schenk, H., Oschlies, A. & Traulsen, A. Immediate action is the best strategy when facing uncertain climate change. *Nature Communications* **9**, 2566 (2018).

[9] Sigmund, K. *The Calculus of Selfishness* (Princeton Univ. Press, Princeton, NJ, 2010).

[10] McNamara, J. M. Towards a richer evolutionary game theory. *Journal of The Royal Society Interface* **10**, 20130544 (2013).

[11] Bergstrom, T., Blume, L. & Varian, H. On the private provision of public goods. *Journal of Public Economics* **29**, 25–49 (1986).

[12] Diekmann, A. Volunteer's dilemma. *Journal of Conflict Resolution* **29**, 605–610 (1985).

[13] Weesie, J. Asymmetry and timing in the volunteer's dilemma. *Journal of Conflict Resolution* **37**, 569–590 (1993).

[14] Vasconcelos, V. V., Santos, F. C., Pacheco, J. M. & Levin, S. A. Climate policies under wealth inequality. *Proceedings of the National Academy of Sciences USA* **111** (2014).

[15] Abou Chakra, M. & Traulsen, A. Under high stakes and uncertainty the rich should lend the poor a helping hand. *Journal of Theoretical Biology* **341**, 123–130 (2014).

[16] Akin, E. What you gotta know to play good in the iterated prisoner's dilemma. *Games* **6**, 175–190 (2015).

[17] Press, W. H. & Dyson, F. D. Iterated prisoner's dilemma contains strategies that dominate any evolutionary opponent. *PNAS* **109**, 10409–10413 (2012).

[18] Akin, E. The iterated prisoner's dilemma: Good strategies and their dynamics. In Assani, I. (ed.) *Ergodic Theory, Advances in Dynamics*, 77–107 (de Gruyter, Berlin, 2016).

[19] Stewart, A. J. & Plotkin, J. B. Collapse of cooperation in evolving games. *Proceedings of the National Academy of Sciences USA* **111**, 17558 – 17563 (2014).

[20] Szolnoki, A. & Perc, M. Defection and extortion as unexpected catalysts of unconditional cooperation in structured populations. *Scientific Reports* **4**, 5496 (2014).

[21] Hilbe, C., Traulsen, A. & Sigmund, K. Partners or rivals? Strategies for the iterated prisoner's dilemma. *Games and Economic Behavior* **92**, 41–52 (2015).

[22] Hilbe, C., Chatterjee, K. & Nowak, M. A. Partners and rivals in direct reciprocity. *Nature Human Behaviour* **2**, 469–477 (2018).

[23] Ledyard, J. O. Public goods: A survey of experimental research. In Kagel, J. H. & Roth, A. E. (eds.) *The Handbook of Experimental Economics* (Princeton Univ. Press, 1995).

[24] Zelmer, J. Linear public goods experiments: A meta-analysis. *Experimental Economics* **6**, 199–310 (2003).

[25] Buckley, E. & Croson, R. Income and wealth inequality in the voluntary provision of linear public goods. *Journal of Public Economics* **90**, 935–955 (2006).

[26] Hargreaves Heap, S. P., Ramalingam, A. & Stoddard, B. Endowment inequality in public good games: A re-examination. *Economics letters* **146**, 4–7 (2016).

[27] Cherry, T. L., Kroll, S. & Shogren, J. F. The impact of endowment heterogeneity and origin on public good contributions: Evidence from the lab. *Journal of Economic Behavior & Organization* **57**, 357–365 (2005).

[28] Kesternich, M., Lange, A. & Sturm, B. On the performance of rule-based contribution schemes under endowment heterogeneity. *Experimental Economics* **21**, 180–204 (2018).

[29] Gächter, S., Mengel, F., Tsakas, E. & Vostroknutov, A. Growth and inequality in public good provision. *Journal of Public Economics* **150**, 1–13 (2017).

[30] Cadigan, J., Wayland, P. T., Schmitt, P. & Swope, K. An experimental dynamic public goods game with carryover. *Journal of Economic Behavior & Organization* **80**, 523–531 (2011).

[31] Nishi, A., Shirado, H., Rand, D. G. & Christakis, N. A. Inequality and visibility of wealth in experimental social networks. *Nature* **526**, 426–429 (2015).

[32] Gächter, S. Visible inequality breeds more inequality. *Nature* **526**, 333–334 (2015).

[33] Chan, K. S., Mestelman, S., Moir, R. & Muller, R. A. Heterogeneity and the voluntary provision of public goods. *Experimental Economics* **2**, 5–30 (1999).

[34] Tavoni, A., Dannenberg, A., Kallis, G. & Löschel, A. Inequality, communication, and the avoidance of disastrous climate change in a public goods game. *Proceedings of the National Academy of Sciences USA* **108**, 11825–11829 (2011).

[35] Margreiter, M., Sutter, M. & Dittrich, D. Individual and collective choice and voting in common pool resource problems with heterogeneous actors. *Environmental and Resource Economics* **32**, 241–271 (2005).

[36] Noussair, C. & Tan, F. Voting on punishment systems within a heterogeneous group. *Journal of Public Economic Theory* **13**, 661–693 (2011).

[37] Gallier, C., Kesternich, M. & Sturm, B. Voting for burden sharing rules in public goods games. *Environmental and Resource Economics* **67**, 535–557 (2017).

[38] Keser, C., Markstädter, A. & Schmidt, M. Mandatory minimum contributions, heterogeneous endowments and voluntary public-good provision. *Games and Economic Behavior* **101**, 291–310 (2017).

[39] Reuben, E. & Riedl, A. Enforcement of contribution norms in public good games with heterogeneous populations. *Games and Economic Behavior* **77**, 122–137 (2013).

[40] Mao, A., Dworkin, L., Suri, S. & Watts, D. J. Resilient cooperators stabilize long-run cooperation in the finitely repeated prisoner's dilemma. *Nature Communications* **8**, 13800 (2017).

[41] von Gerwen, N., Buskens, V. & van der Lippe, T. Individual training and employees' cooperative behavior: Evidence from a contextualized laboratory experiment. *Rationality and Society* **30**, 432–462 (2018).

[42] Mailath, G. J. & Samuelson, L. *Repeated games and reputations* (Oxford Univ. Press, Oxford, UK, 2006).

[43] Selten, R. Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory* **4**, 25–55 (1975).

[44] Fudenberg, D. & Tirole, J. *Game Theory* (MIT Press, Cambridge, 1998), 6th edn.

[45] Traulsen, A., Pacheco, J. M. & Nowak, M. A. Pairwise comparison and selection temperature in evolutionary game dynamics. *Journal of Theoretical Biology* **246**, 522–529 (2007).

[46] Nowak, M. A., Sasaki, A., Taylor, C. & Fudenberg, D. Emergence of cooperation and evolutionary stability in finite populations. *Nature* **428**, 646–650 (2004).

[47] Hofbauer, J. & Sigmund, K. *Evolutionary Games and Population Dynamics* (Cambridge University Press, Cambridge, UK, 1998).

[48] Ohtsuki, H. Stochastic evolutionary dynamics of bimatrix games. *Journal of Theoretical Biology* **264**, 136–142 (2010).

[49] Nowak, M. A. & Sigmund, K. A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner's Dilemma game. *Nature* **364**, 56–58 (1993).

[50] Hauert, C. & Schuster, H. G. Effects of increasing the number of players and memory size in the iterated prisoner's dilemma: a numerical approach. *Proceedings of the Royal Society B* **264**, 513–519 (1997).

[51] Pinheiro, F. L., Vasconcelos, V. V., Santos, F. C. & Pacheco, J. M. Evolution of all-or-none strategies in repeated public goods dilemmas. *PLoS Comput Biol* **10**, e1003945 (2014).

[52] Hilbe, C., Wu, B., Traulsen, A. & Nowak, M. A. Cooperation and control in multiplayer social dilemmas. *Proceedings of the National Academy of Sciences USA* **111**, 16425–16430 (2014).

[53] Maynard Smith, J. *Evolution and the Theory of Games* (Cambridge University Press, Cambridge, 1982).

[54] Selten, R. & Hammerstein, P. Gaps in harley's argument on evolutionarily stable learning rules and in the logic of "Tit for Tat". *Behavioral and Brain Sciences* **7**, 115–116 (1984).

[55] Boyd, R. & Lorberbaum, J. No pure strategy is evolutionary stable in the iterated prisoner's dilemma game. *Nature* **327**, 58–59 (1987).

[56] Farrell, J. & Ware, R. Evolutionary stability in the repeated prisoner's dilemma. *Theoretical Population Biology* **36**, 161–166 (1989).

[57] Lorberbaum, J. M. D. No strategy is evolutionary stable in the repeated Prisoner's Dilemma. *Journal of Theoretical Biology* **168**, 117–130 (1994).

[58] Bendor, J. & Swistak, P. Types of evolutionary stability and the problem of cooperation. *Proceedings of the National Academy of Sciences USA* **92**, 3596–3600 (1995).

[59] van Veelen, M., García, J., Rand, D. G. & Nowak, M. A. Direct reciprocity in structured populations. *Proceedings of the National Academy of Sciences USA* **109**, 9929–9934 (2012).

[60] García, J. & van Veelen, M. In and out of equilibrium I: Evolution of strategies in repeated games with discounting. *Journal of Economic Theory* **161**, 161–189 (2016).

[61] García, J. & van Veelen, M. No strategy can win in the repeated prisoner's dilemma: Linking game theory and computer simulations. *Frontiers in Robotics and AI* **5**, 102 (2018).

[62] Boyd, R. Mistakes allow evolutionary stability in the repeated Prisoner's Dilemma game. *Journal of Theoretical Biology* **136**, 47–56 (1989).

[63] Fudenberg, D. & Maskin, E. Evolution and cooperation in noisy repeated games. *American Economic Review* **80**, 274–279 (1990).

[64] Lindgren, K. Evolutionary dynamics in game-theoretic models. In Arthur, W. B., Durlauf, S. N. & Lane, D. A. (eds.) *The Economy as an Evolving Complex System II*, 337–368 (Addison-Wesley, Reading MA, 1997).

[65] Hilbe, C., Martinez-Vaquero, L. A., Chatterjee, K. & Nowak, M. A. Memory-n strategies of direct reciprocity. *Proceedings of the National Academy of Sciences USA* **114**, 4715–4720 (2017).

[66] García, J. & Traulsen, A. The structure of mutations and the evolution of cooperation. *PLoS One* **7**, e35287 (2012).

[67] Stewart, A. J. & Plotkin, J. B. The evolvability of cooperation under local and non-local mutations. *Games* **6**, 231–250 (2015).

[68] Axelrod, R. *The evolution of cooperation* (Basic Books, New York, NY, 1984).

[69] Paolacci, G., Chandler, J. & Ipeirotis, P. G. Running experiments on Amazon Mechanical Turk. *Judgment and Decision Making* **5**, 411–419 (2010).

[70] Buhrmester, M., Kwang, T. & Gosling, S. D. Amazon's Mechanical Turk: A new source of inexpensive, yet high-quality data? *Perspectives on Psychological Science* **6**, 3–5 (2011).

[71] Horton, J. J., Rand, D. G. & Zeckhauser, R. J. The online laboratory: Conducting experiments in a real labor market. *Experimental Economics* **4**, 399–425 (2011).

[72] Hendriks, A. SoPHIE – Software Platform for Human Interaction Experiments. *Working paper* (2012).

[73] Hauser, O. P., Hendriks, A., Rand, D. G. & Nowak, M. A. Think global, act local: Preserving the global commons. *Scientific Reports* **6**, 36079 (2016).

[74] Hauser, O. P., Kraft-Todd, G. T., Rand, D. G., Nowak, M. A. & Norton, M. I. Invisible inequality leads to punishing the poor and rewarding the rich. *Academy of Management Proceedings* **2016**,

13841 (2016).

[75] Heinze, G. & Schemper, M. A solution to the problem of separation in logistic regression. *Statistics in Medicine* **21**, 2409–2419 (2002).

[76] Firth, D. Bias reduction of maximum likelihood estimates. *Biometrika* **80**, 27–38 (1993).

[77] Stekhoven, D. J. & Bühlmann, P. MissForest – non-parametric missing value imputation for mixed-type data. *Bioinformatics* **28**, 112–118 (2012).

[78] Rice, W. R. Analyzing tables of statistical tests. *Evolution* **43**, 223–225 (1989).

[79] Bland, J. M. & Altman, D. G. Multiple significance tests: the Bonferroni method. *BMJ* **310**, 170 (1995).

[80] Hauser, O. P. & Norton, M. I. (Mis)perceptions of inequality. *Current Opinion in Psychology* **18**, 21–25 (2017).

[81] Sigmund, K. Punish or perish? Retaliation and collaboration among humans. *Trends in Ecology and Evolution* **22**, 593–600 (2007).

[82] Gächter, S., Renner, E. & Sefton, M. The long-run benefits of punishment. *Science* **322**, 1510 (2008).

[83] Egas, M. & Riedl, A. The economics of altruistic punishment and the maintenance of cooperation. *Proceedings of the Royal Society B* **275**, 871–878 (2008).

[84] Boyd, R., Gintis, H. & Bowles, S. Coordinated punishment of defectors sustains cooperation and can proliferate when rare. *Science* **328**, 617–620 (2010).

[85] Sigmund, K., De Silva, H., Traulsen, A. & Hauert, C. Social learning promotes institutions for governing the commons. *Nature* **466**, 861–863 (2010).

[86] Traulsen, A., Röhl, T. & Milinski, M. An economic experiment reveals that humans prefer pool punishment to maintain the commons. *Proceedings of the Royal Society B* **279**, 3716–3721 (2012).

[87] Hilbe, C., Traulsen, A., Röhl, T. & Milinski, M. Democratic decisions establish stable authorities that overcome the paradox of second-order punishment. *Proceedings of the National Academy of Sciences USA* **111**, 752–756 (2014).

[88] Abdallah, S. *et al.* Corruption drives the emergence of civil society. *Journal of The Royal Society Interface* **11**, 20131044 (2014).

[89] Muthukrishna, M., Francois, P., Pourahmadi, S. & Henrich, J. Corrupting cooperation and how anti-corruption strategies may backfire. *Nature Human Behaviour* **1**, 0138 (2017).