



Efficiency and resilience of cooperation in asymmetric social dilemmas

Valentin Hübner^a, Manuel Staab^b, Christian Hilbe^c, Krishnendu Chatterjee^{a,1,2}, and Maria Kleshina^{d,1}

Edited by Brian Skyrms, University of California, Irvine, CA; received September 8, 2023; accepted January 17, 2024

Direct reciprocity is a powerful mechanism for cooperation in social dilemmas. The very logic of reciprocity, however, seems to require that individuals are symmetric, and that everyone has the same means to influence each others' payoffs. Yet in many applications, individuals are asymmetric. Herein, we study the effect of asymmetry in linear public good games. Individuals may differ in their endowments (their ability to contribute to a public good) and in their productivities (how effective their contributions are). Given the individuals' productivities, we ask which allocation of endowments is optimal for cooperation. To this end, we consider two notions of optimality. The first notion focuses on the resilience of cooperation. The respective endowment distribution ensures that full cooperation is feasible even under the most adverse conditions. The second notion focuses on efficiency. The corresponding endowment distribution maximizes group welfare. Using analytical methods, we fully characterize these two endowment distributions. This analysis reveals that both optimality notions favor some endowment inequality: More productive players ought to get higher endowments. Yet the two notions disagree on how unequal endowments are supposed to be. A focus on resilience results in less inequality. With additional simulations, we show that the optimal endowment allocation needs to account for both the resilience and the efficiency of cooperation.

social dilemmas | direct reciprocity | inequality | public good games | evolution of cooperation

Everyday life is rich in situations where individuals have to decide whether to act for the benefit of the group or their individual gain. These situations are commonly referred to as social dilemmas (1, 2). For pro-social behavior to be maintained in these settings, it takes some mechanism that enables individuals to overcome selfish interests (3). Direct reciprocity is one such mechanism (4–6). It requires individuals to interact repeatedly, so that previous actions may shape future decisions. As a result, even in the absence of explicit punishments, cooperation can evolve and be stable.

Most previous studies on the evolution of reciprocity focus on fully symmetric interactions (7–26). In these studies, individuals are perfectly interchangeable. This assumption plays a critical role in the context of direct reciprocity, because it implies that the ability to increase or reduce an opponent's payoff is identical across individuals. If, however, individuals differ in their costs and benefits of cooperation, some individuals might be harder to discipline than others, making cooperation more difficult to sustain. This has, for instance, been shown in the context of endowment inequality (27–43). Endowment inequality is often identified with real-world inequalities in income or wealth, which have a negative impact on social outcomes more generally (44–48). Such observations suggest that if individuals are otherwise symmetric, endowment inequality ought to be as small as possible to promote cooperation. However, in addition to endowment inequality, individuals often differ along multiple other dimensions, such as their level of skill. In such a context, recent studies suggest that a perfectly equal endowment distribution may not be optimal for cooperation either (49–51). This raises the question what the optimal level of endowment inequality is. This paper aims to provide an answer to that question.

Taking the framework of Hauser et al. (52) as a starting point, we study repeated linear public good games among asymmetric players. Our baseline model contains two sources of asymmetry. First, players may differ in their endowments, which influences how much they can contribute to the public good. Second, players may differ in their productivities, which influences how effective contributions are. Given the players' productivities, we ask how endowments should be optimally allocated. To tackle that question, we introduce two notions of optimality. First, we characterize the endowment distribution that results in the highest "resilience of cooperation". Here, individuals are able to enforce cooperation even as the game's continuation probability approaches the theoretical

Significance

Many collective-action problems require unrelated individuals to cooperate. Such problems can be difficult to solve, particularly when actors are heterogeneous. For example, actors often differ in their stakes in the game, or in the effectiveness of their contributions. Here, we explore the impact of such heterogeneities on the evolution of reciprocal cooperation in groups of arbitrary size. We first show that there is an optimal degree of (endowment) inequality that maximizes the stability of cooperative outcomes. We also demonstrate that, counter-intuitively, improving the stability of cooperation does not need to maximize social welfare. Additional simulations of a learning process suggest that individuals tend to balance this trade-off between efficiency and stability of cooperation.

Author contributions: V.H., M.S., C.H., K.C., and M.K. designed research; V.H., M.S., and M.K. performed research; V.H. and M.K. analyzed data; and V.H., M.S., C.H., K.C., and M.K. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

Copyright © 2024 the Author(s). Published by PNAS. This article is distributed under Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 (CC BY-NC-ND).

¹K.C. and M.K. contributed equally to this work.

²To whom correspondence may be addressed. Email: krishnendu.chatterjee@ista.ac.at.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2315558121/-/DCSupplemental>.

Published February 26, 2024.

minimum. Second, we characterize the endowment distribution that exhibits the highest “efficiency of cooperation.” This distribution maximizes social welfare in the best possible equilibrium. We identify those two optimal distributions for any form of heterogeneity in individual productivities and any group size.

We find that according to both notions, more productive players ought to get higher endowments. The exact magnitude of this optimal endowment inequality, however, depends on which notion is used, and on the parameters of the game. In particular, we identify scenarios where resilience of cooperation requires endowments to be almost equal even though productivities are not. Conversely, we also describe scenarios in which minor differences in productivities result in major differences in endowments. As a general rule, we find that the efficiency-maximizing endowment distribution is always more unequal than the resilience-maximizing one. This suggests that there is a non-trivial trade-off between the resilience of cooperation and efficiency. To further study this trade-off, we simulate learning dynamics among interacting individuals. These simulations suggest that the endowment distribution that performs best lies on a Pareto frontier between efficiency and resilience. Where exactly that point is located depends on the chosen parameter values. When parameters are generally favorable to cooperation, payoffs are highest when endowments are close to the efficiency-

maximizing distribution. In contrast, in noisy environments in which cooperation is generally difficult to sustain, payoffs are higher when the endowment distribution prioritizes resilience.

This work highlights how different objectives, such as efficiency or resilience, have different implications for the optimal allocation of endowments within groups. While both objectives tolerate some endowment inequality, this inequality needs to be smaller when cooperation is to be resilient.

Results

Model. We consider a repeated linear public good game among n players. At every time step t , each player i receives an endowment e_i . Subsequently, each player decides how much of the endowment to contribute toward the public good, $c_i(t)$, and how much to consume individually, $e_i - c_i(t)$ (Fig. 1A). The contribution of a player i is multiplied by a productivity factor r_i , with $1 < r_i < n$, which may differ across players (Fig. 1B). The total amount of the public good produced equals the sum of all effective contributions $\sum_{i=1}^n r_i c_i(t)$. This amount is equally divided among all players, independently of their contributions (Fig. 1C and D). Individual payoffs $\pi_i(t)$ are determined by the quantity of the public good received, and the individually consumed shares of the endowments (Fig. 1E),

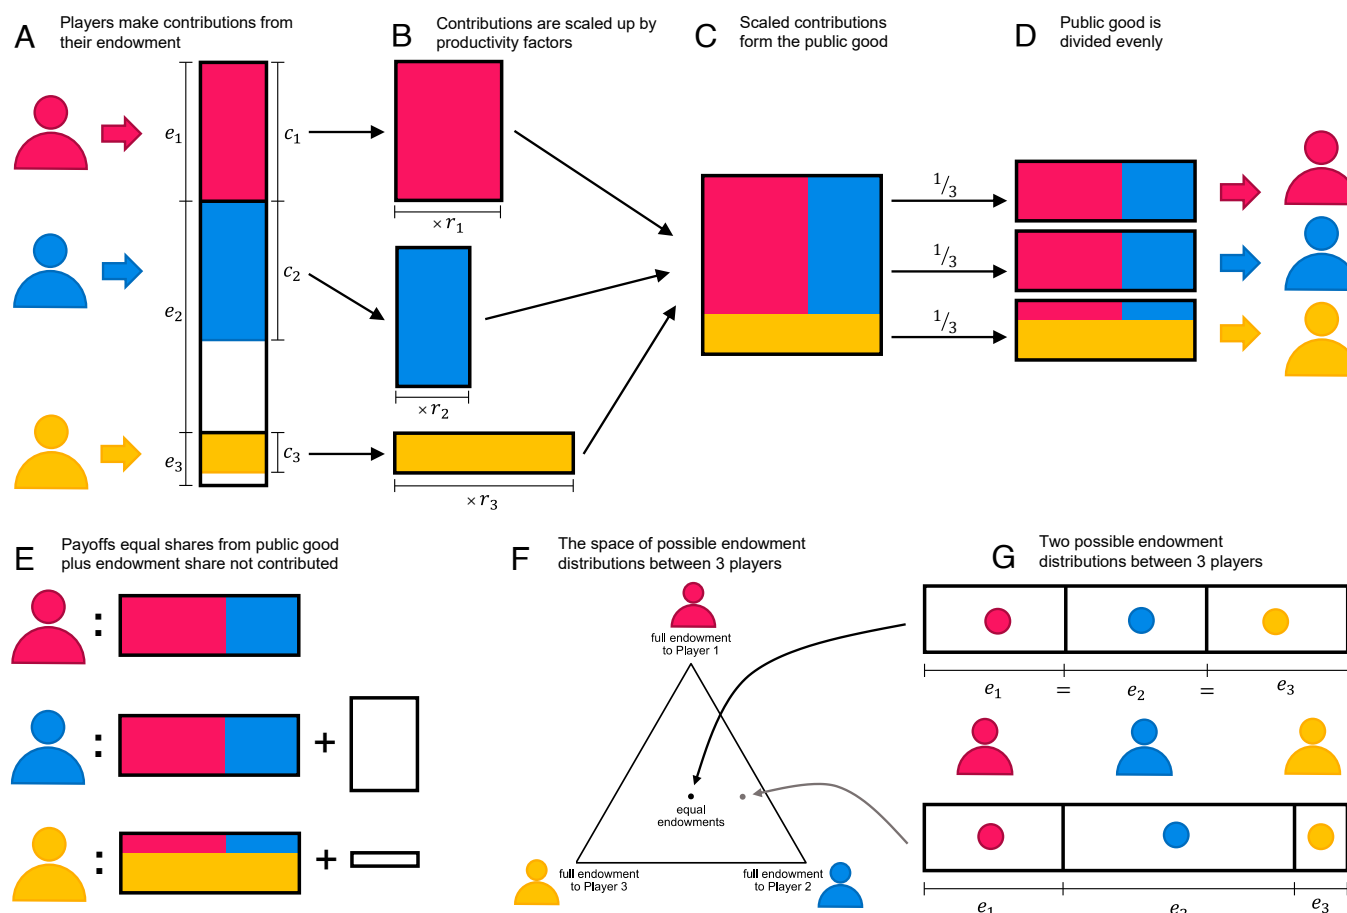


Fig. 1. Schematic representation of the model. Players engage in a repeated linear asymmetric public good game. In every round, each player i receives an endowment e_i . (A) Players choose how much to contribute toward the production of the public good, c_i , from their available endowment, e_i . (B) All individual contributions are multiplied by individual productivity factors, r_i . (C) The size of the public good is defined as the sum of all effective contributions. (D) After its production, the public good is divided equally among all players. (E) Individual payoffs are equal to the n th share of the public good plus the remaining share of the endowment that players did not contribute toward the public good. (F) Without loss of generality, we assume $\sum_{i=1}^n e_i = 1$. In the case of a three-player game, we can represent the endowment distributions in a simplex, where each point corresponds to a vector $\mathbf{e} = (e_1, e_2, e_3)$. (G) We aim to identify the optimal endowment distribution with respect to different objectives.

$$\pi_i(t) = e_i - c_i(t) + \frac{1}{n} \sum_{j=1}^n r_j c_j(t). \quad [1]$$

Without loss of generality, we assume that endowments are normalized such that $\sum_{i=1}^n e_i = 1$. Accordingly, we refer to the vector $\mathbf{e} = (e_1, \dots, e_n)$ as an endowment distribution. This distribution summarizes how endowments are allocated among the players (Fig. 1 *F* and *G*). For example, the vector $\mathbf{e} = (1/n, \dots, 1/n)$ describes an equal allocation.

If the above game is only played for a single round, full defection is the only equilibrium, for any endowment distribution. However, here we assume that after each round t , there is another round with a fixed continuation probability $0 < \delta < 1$. Equivalently, one may also interpret our setup as that of a game with infinitely many rounds, and δ as the extent to which players care about their future payoffs (53). In each case, Player i 's expected payoff over all rounds, with a normalizing factor of $1 - \delta$, is given by

$$\hat{\pi}_i = (1 - \delta) \sum_{t=0}^{\infty} \delta^t \pi_i(t). \quad [2]$$

For repeated games, the Folk theorem (54, 55) states that any individually rational outcome can be sustained in a Subgame Perfect Nash Equilibrium (SPNE), provided δ is sufficiently close to 1. In the following, we are particularly interested in equilibria with full cooperation, meaning that all players choose $c_i(t) = e_i$ in every single round t (By this definition, a player i who happens to get no endowment is considered fully cooperative, even though the player contributes $c_i(t) = e_i = 0$ every round). In general, whether or not an equilibrium with full cooperation exists depends on how endowments are allocated. For large δ , there are generally many different endowment distributions for which a fully cooperative equilibrium exists. But as δ decreases, the respective set of endowment distributions shrinks. In that sense, endowment distributions differ in how resilient cooperation is to adverse circumstances. Moreover, due to the variation in individual productivities, different endowment distributions result in different levels of welfare even if everyone fully cooperates. We therefore investigate the resilience of cooperation and the resulting welfare implications.

More specifically, we distinguish two objectives: First, we aim to find the endowment distribution that sustains full cooperation at the lowest δ . We call this the “resilience-maximizing endowment distribution.” Second, we are interested in the endowment distribution that allows for the highest group welfare in equilibrium. We refer to this as the “efficiency-maximizing endowment distribution.”

The Resilience-Maximizing Endowment Distribution. It turns out that there is a surprisingly elegant characterization of the resilience-maximizing endowment distribution. In the following, we provide a summary of our corresponding findings. All details and formal derivations are described in *SI Appendix*.

First, we show that for an arbitrary (but fixed) set of productivities $\mathbf{r} = (r_1, \dots, r_n)$, there is always an endowment distribution \mathbf{e} and a continuation probability $\delta < 1$ such that full cooperation is possible. Next, if also the values of δ and \mathbf{e} are fixed, we prove that full cooperation is sustainable in an SPNE exactly if

$$(\delta D - I_n)\mathbf{e} \geq \mathbf{0}. \quad [3]$$

Here, D is an $n \times n$ matrix with entries $D_{ij} = \frac{r_j}{n-r_i}$ for $i \neq j$ and $D_{ii} = 0$, and I_n is the $n \times n$ identity matrix. We observe that for the given endowment distribution \mathbf{e} , there exists a minimal continuation probability $\delta_{\min}(\mathbf{e})$ that satisfies Eq. 3. Hence full cooperation is sustainable in a SPNE if and only if $\delta \geq \delta_{\min}(\mathbf{e})$. Because δ can be interpreted as the patience of players, or how much they value their future payoffs, this lower bound on δ can be considered to be a measure of how hard it is to sustain cooperation with the given endowment distribution. The lower this minimum $\delta_{\min}(\mathbf{e})$, the easier it is to sustain cooperation.

Based on this observation, we define the resilience-maximizing endowment distribution \mathbf{e}^* to be the one with the smallest value of $\delta_{\min}(\mathbf{e})$, that is, $\mathbf{e}^* := \arg \min_{\mathbf{e}} \delta_{\min}(\mathbf{e})$. We use the notation $\delta_{\min}^* := \delta_{\min}(\mathbf{e}^*)$ for the corresponding minimal continuation probability. Using inequality Eq. 3, we can derive \mathbf{e}^* and δ_{\min}^* for any number of players n and individual productivities \mathbf{r} . We show that \mathbf{e}^* is exactly the Perron eigenvector of D , and the corresponding eigenvalue is equal to $(\delta_{\min}^*)^{-1}$ (*SI Appendix, Text S3*). This provides a simple method for calculating \mathbf{e}^* for any set of parameters and any group size. For the special case of a two-player game, we recover Hauser et al.'s (27) result that the resilience-maximizing endowment distribution is equal to

$$e_1^* = \frac{\sqrt{r_2(2-r_2)}}{\sqrt{r_1(2-r_1)} + \sqrt{r_2(2-r_2)}}$$

and

$$e_2^* = \frac{\sqrt{r_1(2-r_1)}}{\sqrt{r_1(2-r_1)} + \sqrt{r_2(2-r_2)}}.$$

Based on our characterization for n players, we can derive general properties of the resilience-maximizing endowment distribution \mathbf{e}^* . We find that the relationship between e_i^* and r_i is always order preserving. That is, more productive players always need to have a larger endowment than less productive players in order to guarantee the highest resilience of cooperation. Nonetheless, the degree of endowment inequality according to \mathbf{e}^* may vary significantly. It depends on the ratio between players' individual productivities and the size of the group, r_i/n . The smaller the productivities are in relation to n , the more equal \mathbf{e}^* is. In particular, fixing individual productivities at some level and increasing n results in \mathbf{e}^* getting arbitrarily close to $(1/n, \dots, 1/n)$. To see this, we note that for large n , the off-diagonal entries of the matrix D approach $D_{ij} = r_j/n$, which is independent of i . For the resulting matrix D , the uniform distribution is a right eigenvector.

We illustrate this effect in Fig. 2. When the players' productivities are comparably low (Fig. 2*A*), we observe a resilience-maximizing endowment allocation that is approximately uniform (Fig. 2*B*). However, due to the low productivities, cooperation is not very resilient (Fig. 2*C*). Keeping the variance in productivities fixed while increasing their overall level (such that Player 1's productivity is almost equal to n , Fig. 2*D*) results in a very unequal endowment allocation. Now, Player 1 receives almost all of the endowment (Fig. 2*E*). Yet, despite the high inequality in endowments, cooperation becomes more resilient (Fig. 2*F*). We further extend our argument by doubling the number of players, while keeping the productivities fixed (Fig. 2*G*). This reduces r_i/n by half. This new six-player game with identical variance in productivities again results in an almost uniform distribution \mathbf{e}^* (Fig. 2*H*). The resilience of cooperation is intermediate in this case.

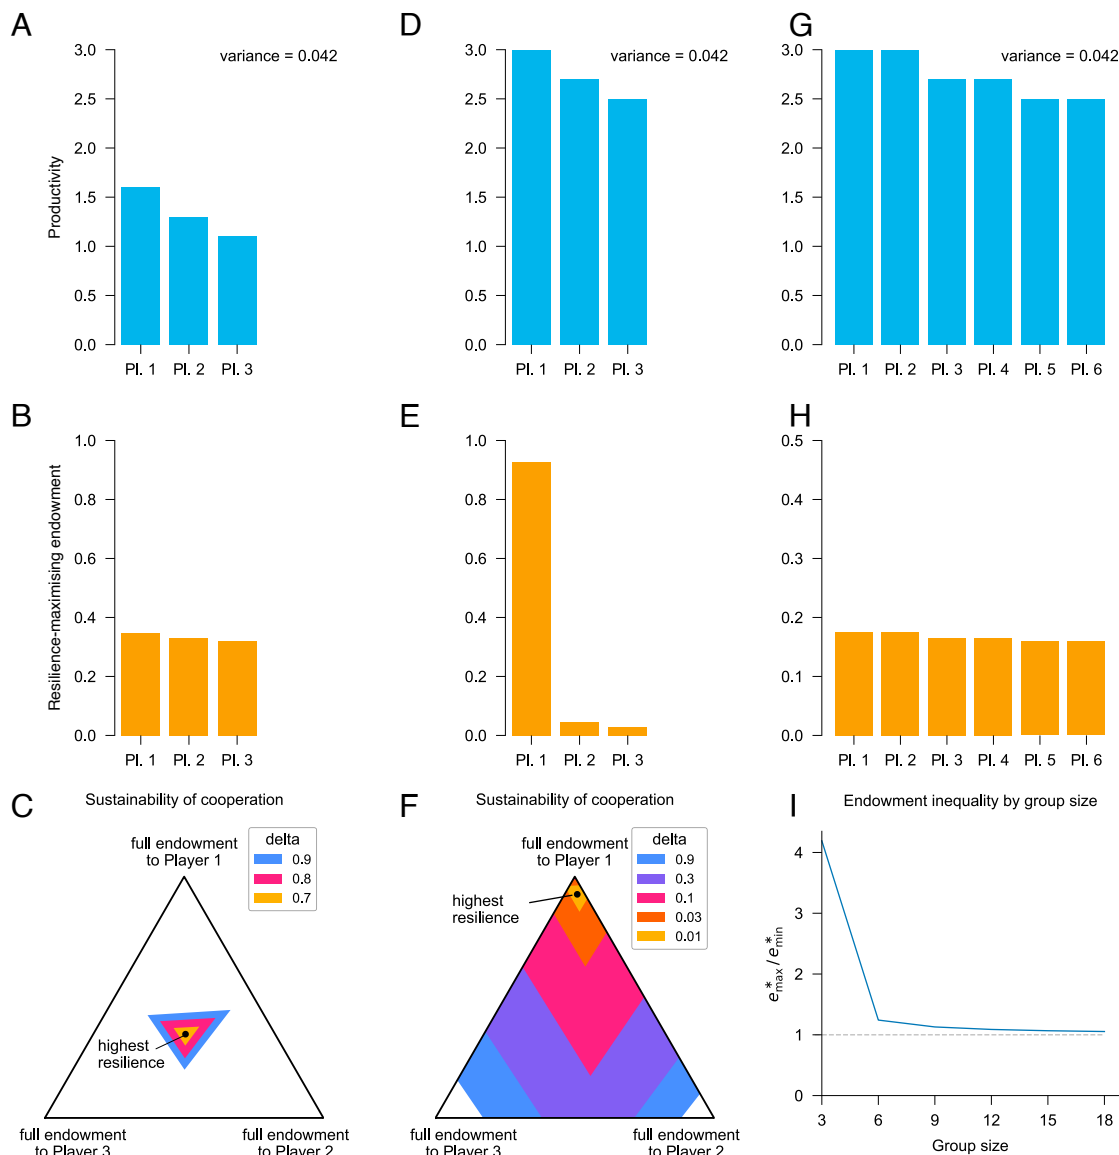


Fig. 2. Resilience-maximizing endowment distribution. (A) To demonstrate the degree of inequality in the resilience-maximizing endowment distribution, we construct three examples of games, all of which have the same level of heterogeneity in productivities. In this first example, three players with comparatively low productivities interact. (B) When the ratio r_i/n is low, then the resilience-maximizing endowment distribution is close to $(1/n, \dots, 1/n)$. (C) With productivities close to 1, cooperation is challenging: $\delta_{\min} = 0.620$. (D–F) Cooperation becomes more attractive and much easier to sustain. $\delta_{\min} = 0.005$. The resulting degree of inequality of e^* increases. (G–I) With high productivities, but more players, the endowment distribution is again close to $(1/n, \dots, 1/n)$. The required continuation probability is now $\delta_{\min} = 0.238$, which is lower than in the previous example, but higher than in the example with low productivities. (I) We demonstrate the general principle by systematically varying group size while keeping productivities fixed (up to multiplicity). We plot the degree of inequality of the resulting resilience-maximizing endowment distribution measured by the ratio between the highest and the lowest endowments in the allocation. As the ratio r_i/n decreases, the resilience-maximizing endowment allocation approaches $(1/n, \dots, 1/n)$.

We can formalize this result by deriving an upper bound on the relative difference between players' endowments with respect to e^* . This upper bound is given by

$$\frac{\max_i e_i^*}{\min_i e_i^*} \leq \frac{n-1}{n - \max_i r_i}. \quad [4]$$

In particular, if there is some number k such that productivities do not exceed n/k , then the absolute difference in endowments is bounded by $(\max_i e_i^* - \min_i e_i^*) \leq 1/(k-1)$.

Efficiency-Maximizing Endowment Distribution. The social-dilemma nature of the game implies that higher cooperation achieves greater group welfare. However, not all endowment

allocations allow for full contributions in equilibrium and, due to the individual heterogeneity in productivity, full cooperation yields different levels of welfare with different endowment distributions.

We define welfare as the sum of the individual (expected) payoffs. If all players contribute fully, the group welfare can be expressed as a function of endowments:

$$\Phi(\mathbf{e}) := \sum_{i=1}^n \hat{\pi}_i = \sum_{i=1}^n r_i e_i, \quad [5]$$

where $\hat{\pi}_i$ is as defined in equation Eq. 2.

Maximization of the group's welfare constitutes an optimization problem of finding an endowment distribution \mathbf{e}^\dagger under




Fig. 3. Trade-off between efficiency and resilience of cooperation. We demonstrate the difference between the resilience- and efficiency-maximizing endowment distributions in a three-player example. (A) We chose the productivity vector to be $r = (2.7, 1.5, 1.1)$. (B) The resulting resilience-maximizing endowment distribution yields the total social welfare of $\Phi = 2.025$. (C) We find that the efficiency-maximizing endowment at $\delta = 0.3$ is more unequal and yields a total group payoff of $\Phi = 2.174$. (D and E) We formulate a multi-objective optimization problem where both the resilience and efficiency are varied. The Pareto optimal values are shown by the pink line. (F) We run simulations to test which of the endowment distributions performs best when players adopt straz tegies based on a stochastic learning process. We find that in general, the highest cooperation levels are achieved along the Pareto frontier. Indeed, the total maximum group payoff of 1.729 is achieved at $e_{\max W} = (0.65, 0.35, 0)$.

which full cooperation is sustainable and which maximizes endowment distribution. We refer to e^\dagger as an “efficiency-maximizing endowment distribution.” While finding an explicit expression for e^\dagger is not possible in general (SI Appendix, Section S4), we can obtain numerically exact solutions for any group size n . Furthermore, we can fully characterize the general functional form of e^\dagger in the two-player case (considering without loss of generality $r_1 \geq r_2$) as

$$e_1^\dagger = \frac{\delta r_2}{2 - r_1 + \delta r_2} \quad \text{and} \quad e_2^\dagger = \frac{2 - r_1}{2 - r_1 + \delta r_2}. \quad [6]$$

As expected, the efficiency-maximizing endowment distribution allocates larger shares of the endowment to more productive players. While this effect is similar to the effect of the resilience-maximizing endowment distribution, the resulting degree of inequality differs (Fig. 3 A–C). In fact, there exist parameters r and δ such that the efficiency-maximizing distribution results in an exclusion of the least productive players by allocating them a zero-share of the total endowment. This is in stark contrast to the resilience-maximizing endowment distribution; for e^* , we prove that all players are always allocated a positive share (SI Appendix, Corollary 7).

Trade-Off between Efficiency and Resilience. Since the endowment distributions e^* and e^\dagger are generally not the same, we further analyze the relation between them. There are two possible cases: First, the resilience-maximizing endowment distribution e^* simultaneously also achieves maximal efficiency. This occurs exactly if $\delta = \delta_{\min}^*$ (in which case there is a unique endowment distribution that can sustain cooperation), or if all players have the same productivity (in which case every endowment distribution has the same efficiency). Second, in all other cases, we can prove that under any measure of inequality, the efficiency-maximizing endowment distribution e^\dagger is always more unequal than the resilience-maximizing distribution e^* .

Given there is a trade-off between the two objectives in most settings, we combine them into a multi-objective optimization setup. We visualize the resulting Pareto frontier between the resilience of cooperation and its efficiency in Fig. 3 D and E. The pink line indicates the maximum welfare that can be sustained for any value of $\delta_{\min}(e)$ (Fig. 3D). Each point on this line corresponds to an endowment distribution e with the corresponding values of $\delta_{\min}(e)$ and $\Phi(e)$ (Fig. 3E). Generally, the most efficient endowment allocation e^\dagger is located at the boundary of the set of all endowment allocations that allow for full cooperation. Hence, while securing a maximal group payoff, it also poses the greatest strain on the resilience of

cooperation. On the other hand, the resilience-maximizing endowment allocation always requires higher equality while yielding lower welfare, signifying the trade-off between efficiency and resilience.

Dynamics of Cooperation. To complement these static equilibrium results, in the following, we explore when cooperation can emerge for a given set of endowments and productivities. To this end, we no longer assume that players act optimally from the outset. Rather, they adapt their strategies over time to optimize their payoffs. To model this, we use introspection dynamics (27, 56), a learning process where players are repeatedly selected at random to revise their strategies. When selected, players compare their current payoff with the payoff they could have obtained with a randomly generated alternative strategy. The higher the payoff of the alternative, the more likely players switch (as described in detail in *Materials and Methods*). In line with the literature on direct reciprocity (7–26), we assume that individuals can only adopt pure memory-one strategies. That is, players condition their actions only on the outcome of the previous round (6). Moreover, in any given round, players either contribute their entire endowment or nothing at all. The resulting learning dynamics can be represented by a Markov chain. By computing its invariant distribution, we can infer the frequencies of each of the memory-one strategies in the long run.

To start with, we explore the simplest possible case of a game with two players with equal productivities. Here, the resilience-maximizing and efficiency-maximizing endowment distributions coincide at $e_1 = e_2 = 0.5$. In agreement with this prediction, we find that equal endowments are most favorable to the evolution of cooperation (Fig. 4B). We also find, as expected,

that higher values of the selection strength parameter allow for more cooperation (Fig. 4C). The effect of the error rate ϵ however is not monotonous (Fig. 4A). As has been documented in the past, a moderate amount of errors can be beneficial, because errors prevent the neutral invasion of conditionally cooperative strategies like “Win-Stay Lose-Shift” by unconditional cooperators (57). Excessive errors, however, are always detrimental to cooperation, because they render conditionally cooperative strategies unstable.

Next, we look at a scenario with heterogeneity in individual productivities. We find that the endowment distribution that achieves the highest group welfare is located somewhere between the resilience-maximizing and efficiency-maximizing endowment distributions (Fig. 4E). Its exact location depends on the error rate and the selection strength. We observe that an increase in the selection strength results in a clear shift toward higher efficiency of the endowment distribution (Fig. 4F). In contrast, the error rate can have varying effects (Fig. 4E). As a rule of thumb, in more noisy settings (either because of a high error rate, or a low selection strength), allocations close to the resilience-maximizing endowment distribution tend to result in a higher welfare.

In addition, *SI Appendix, Fig. S1* reports the resulting average cooperation rates. As can be seen, endowment allocations for which we observe the maximum group welfare are different from the endowment allocations that maximize cooperation (Fig. 4E and *SI Appendix, Fig. S1E*). To gain intuition for why this is the case, we look at the distribution of strategies for each endowment distribution of interest. Apart from the efficiency- and resilience-maximizing endowment distributions, we also include the endowment distributions where we observe maximum group payoffs, $e_{\max W}$, and maximum cooperation rates, $e_{\max C}$, (*SI Appendix, Fig. S2*), respectively. It appears that

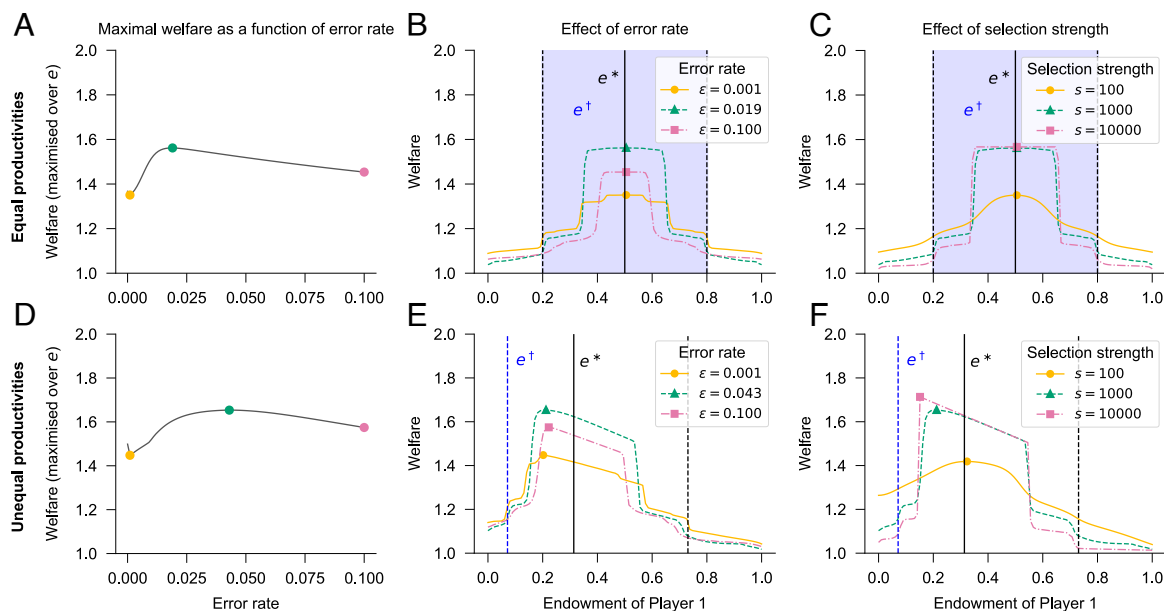


Fig. 4. Evolutionary simulations of the group welfare. In order to gain a deeper insight into the behavior of the dynamics, we provide results of extensive simulations of a two-player game for a wide range of parameters of the dynamics, that is, the error rate and the intensity of selection. (A–C) We first study the evolution of cooperation with equal productivities. Here, the dashed vertical lines bound the region where cooperation is sustainable at $\delta = 1$ according to the analytical model. We choose three error rate values for comparison: very rare errors, very frequent errors, and the error rate $\epsilon^* = 0.019$ that yields the highest welfare at $e = e^*$ (indicated by the solid vertical line). As can be seen, some (rare) errors can help the evolution of cooperation by ensuring stability of cooperative strategies such as WLS (57). Here, there is no unique e^+ , since all endowment distributions where full cooperation is sustainable (shaded in blue), yield identical welfare. Near the boundaries of this interval, we observe very low cooperation rates, while the highest group welfare is observed at the resilience-maximizing endowment in the center. (D–F) Next, we consider a two-player game with unequal productivities given as $r_1 = 1.3$ and $r_2 = 1.9$. We employ the same logic for the choice of the parameters and obtain $\epsilon^* = 0.043$. As can be seen in panel (E), the highest group welfare is no longer attained at e^* but at a point in between e^+ and e^* . For ϵ^* , we denote that point as $e_{\max W}$. It is equal to (0.21, 0.79). As with equal productivities, we find that higher selection strength increases welfare. Here, it also shifts the endowment $e_{\max W}$ closer to e^+ .

WSLS is the most abundant strategy for all of these endowment distributions apart from \mathbf{e}^\dagger , where it is not evolutionarily stable. WSLS is less prone to invasion at \mathbf{e}^* and $\mathbf{e}_{\max C}$ for higher error rates (SI Appendix, Fig. S3).

We also explore the dynamics of cooperation in a three-player game (Fig. 3F). As predicted by the model, cooperation is higher when the most productive player obtains the largest endowment. Similarly, we observe that the endowment distributions with the highest group payoffs all lie close to the Pareto frontier between resilience and welfare, which connects the resilience-maximizing and efficiency-maximizing endowment distribution.

Discussion

Numerous studies have shown that wealth inequality can pose a challenge to cooperation (27–43). Since wealth inequality is abundant in many social settings, addressing it is often an important objective for policy makers. Perhaps one of the most straightforward ways of reducing inequality is through wealth redistribution (58–64). As a naive generalization of the theoretical and experimental findings, one could reach the conclusion that any wealth inequality is detrimental to cooperation and welfare, implying that redistribution should aim for equal allocations. However, our analysis shows that finding an allocation that most easily facilitates cooperation while maximizing welfare is non-trivial.

We show that when productivity differs across people, an equal endowment allocation is neither optimal for the resilience of cooperation nor for the maximization of welfare. Yet, consistent with previous findings (27–43), we also find that excessive levels of endowment inequality cause cooperation to break down. We show how the optimal degree of endowment inequality varies with several parameters. In particular, it depends on the ratio between group size and individual productivities: if the players' productivities are fixed, larger groups require more equal endowments to maximize the resilience of cooperation, all the way to perfect equality in the limit of large group size (Fig. 2).

Our findings also point to a connection between the “resilience of cooperation” and “resilience of biological systems” (?). A higher resilience in our model means that full cooperation can be sustained for a wider range of continuation probabilities. If δ is seen as a parameter of the environment, then an endowment distribution with higher resilience allows for cooperation in environments that are less favorable to cooperation in the sense that they have a lower δ . Interpreted more loosely, when endowments are more resilient, cooperation can withstand greater perturbations of the environment.

To what degree resilience can be sacrificed for efficiency gains depends on the context. We explore this trade-off with introspection dynamics (56). We find that the endowment allocation that generates the highest welfare under these learning dynamics is located on the Pareto frontier between resilience and efficiency, balancing these two objectives. These observations are in line with a behavioral experiment conducted by Hauser et al. (52). Interestingly, $\mathbf{e}_{\max W}$ and $\mathbf{e}_{\max C}$ are close to the endowments chosen for the treatments in the experiment. Similar to our numerical results, the authors find that cooperation rates for both of these allocations are roughly the same (approximately 73%) with a higher welfare achieved at the endowment allocation closer to $\mathbf{e}_{\max W}$.

Throughout this main text, we have focused on full cooperation in linear public good games with asymmetry in endowments and productivities. However, our theoretical results, as presented

in SI Appendix, are valid within a framework that is significantly more general in two aspects. First, we allow for the public good to be distributed unequally (SI Appendix, Text S1). This weakens the public-good character of the game, but logically strengthens our results. Second, we consider arbitrary levels of contributions and derive all results with no restrictions on stochasticity, memory, or time dependence of strategies (SI Appendix, Text S3).

We believe our work makes at least two important contributions to the literature on the evolution of cooperation through direct reciprocity. First, we considerably extend earlier results by Hauser et al. (27). Although they discuss which endowment distributions might maximize resilience (there: “endowment distribution most conducive to cooperation”), their analysis is restricted to groups of size two. Instead, here we provide an elegant formalism that allows us to compute \mathbf{e}^* for any number of players. In this way, we can analyze the interaction between parameters of the game and the optimal degree of inequality. Our method can be used to further study the effects of inequality in more general settings, for example, in structured populations or when allowing for communication or signaling among players. Second, we study the effect of inequality on the interplay between the resilience and efficiency of cooperation. Our results for a general n -player linear public good game indicate that there exist non-trivial trade-offs, which need to be accounted for when deciding on the allocation of wealth. We explore these trade-offs using evolutionary simulations. Overall, our results suggest that a positive degree of inequality can be beneficial for cooperation, in particular in small-group interactions, while in other settings, almost perfect equality is optimal even in the face of intrinsic differences between individuals.

Materials and Methods

Model. Consider a game with n players who interact in an infinite sequence of rounds $t = 0, 1, 2, \dots$. Each player has a fixed positive endowment $e_i \geq 0$, where $\sum_i e_i = 1$ without loss of generality. In each round t , each player i chooses a contribution $c_i(t) \in [0, e_i]$ to make toward the public good. The productivity matrix R , a parameter of the game subject to below constraints, governs the relationship between contributions and payoffs as

$$\boldsymbol{\pi}(t) = \mathbf{e} - \mathbf{c}(t) + R\mathbf{c}(t), \text{ for all } t.$$

There are three constraints on R . First, $R_{ij} \geq 0$ for all i, j , ensuring that an increase in one player's contribution does not decrease any player's return from the public good. Second, $R_{ij} < 1$ for all i , so that a player's one-round payoff is higher the less they contribute. Third, $\sum_j R_{ji} > 1$ for all i , meaning that the group payoff of all players taken together is higher the more Player i contributes. The second and third conditions create a tension between the individual and the collective interest, which makes this game a social dilemma. The restricted model discussed in the main text is the special case where R is of the form $R_{ij} = r_j/n$ for some $\mathbf{r} = (r_1, \dots, r_n)$.

Equilibrium Analysis. We determine when it is possible for a given contribution sequence $(\mathbf{c}(t))_t$ to occur in a SPNE. We introduce a normal form of the productivity matrix R , which we call the zero-diagonal form, and denote it by D . It is given by $D_{ij} = R_{ij}/(1 - R_{ij})$ for all $i \neq j$ and $D_{ii} = 0$ for all i . The game defined by D is equivalent to the game defined by R in the sense that the two games permit exactly the same equilibria, while the fact that the diagonal entries of D are zero simplifies the analysis. We show that a contribution sequence $(\mathbf{c}(t))_t$ is sustainable in a SPNE exactly if $\bar{\mathbf{c}}(t) \leq \delta D\bar{\mathbf{c}}(t + 1)$ for all t (SI Appendix, Text S2 and Theorem 1), where $\bar{\mathbf{c}}(t) = (1 - \delta) \sum_{\tau=0}^{\infty} \delta^\tau \mathbf{c}(t + \tau)$ is the continuation contribution after round t . For a given contribution sequence $(\mathbf{c}(t))_t$, we define $\delta_{\min}((\mathbf{c}(t))_t)$ as the smallest continuation probability δ for which the sequence is sustainable.

Evolutionary Analysis. We study introspection dynamics, a simple learning process (56). Players use pure memory-one strategies where in a given round they either contribute the entire endowment or nothing, that is, $c_i(t) \in \{0, e_i\}$. We represent strategies by a vector $\mathbf{p} = (\rho_{\mathbf{c}})_{\mathbf{c}}$, where \mathbf{c} ranges over the 2^n possible outcomes of one round. Each component $\rho_{\mathbf{c}} \in \{0, 1\}$ specifies whether after a round with outcome \mathbf{c} , a player with a strategy \mathbf{p} contributes their entire endowment ($\rho_{\mathbf{c}} = 1$) or nothing ($\rho_{\mathbf{c}} = 0$). However, in our simulations players are also prone to making errors with a probability $\epsilon > 0$, meaning that they sometimes play an action not prescribed by their strategy. Since that makes the process ergodic and we only focus on its asymptotic behavior, our representation of strategies does not contain an initial move.

At every time step of the learning process, a player is chosen to consider switching their strategy. The player compares their current strategy to a randomly generated, alternative strategy in terms of average payoff and adopts the alternative with probability

$$\rho = \frac{1}{1 + e^{s(\pi_{\text{alt}} - \pi_{\text{cur}})'}}$$

where π_{cur} and π_{alt} are the payoffs of the current strategy and the alternative, respectively. The parameter $s \geq 0$ reflects the strength of selection. Higher values of s correspond to stronger selection.

We use different implementations of the learning process depending on n . For $n = 2$, we calculate the asymptotic distribution by numerically representing the stochastic process as a transition matrix of a Markov chain. We then calculate the average payoffs as expected values under the invariant distribution as

$$\pi_i = \sum_{\mathbf{c}} v_{\mathbf{c}} \cdot \pi_i(\mathbf{c}),$$

where π_i is Player i 's average payoff, \mathbf{v} is the invariant distribution of the Markov chain, and $\pi_i(\mathbf{c})$ is Player i 's payoff in a round with contribution vector \mathbf{c} . For $n = 3$, we run an agent-based simulation for $N = 10^6$ generations and report the average group payoffs.

Parameters used for Figures. Fig. 2 A–C shows a three-player game with productivities $r_1 = 1.6$, $r_2 = 1.3$, $r_3 = 1.1$. In Fig. 2 D–F, the relative differences in productivities are identical, but all values are higher by 1.399: $r_1 = 2.999$, $r_2 = 2.699$, $r_3 = 2.499$. Fig. 2 G and H shows a game with 6 players. The productivity values are distributed as in Fig. 2 D–F, in the sense that for each player in the three-player game, there are two identical players in the six-player game: $r_1 = r_2 = 2.999$, $r_3 = r_4 = 2.699$, $r_5 = r_6 = 2.499$. In

Fig. 2I, we show the value of $e_{\text{max}}^*/e_{\text{min}}^*$ as a measure of inequality for groups of various sizes. In each case, one third of players has productivity 2.7, one third has productivity 1.5, and one third has productivity 1.1.

Fig. 3 analyzes a three-player game with productivities $r_1 = 2.7$, $r_2 = 1.5$, $r_3 = 1.1$. In Fig. 3C, we report the efficiency-maximizing endowment distribution for $\delta = 0.3$. In Fig. 3F, we report evolutionary simulations for the same three-player game with the selection strength $s = 1,000$.

Fig. 4 presents data from evolutionary simulations of a two-player game. Here, players can either contribute their entire endowment e_i or defect by contributing 0 to the public good. In Fig. 4 A–C we report results for a symmetric two-player game with productivities $r_1 = r_2 = 1.6$. Fig. 4A depicts the maximal welfare for each value of the error rate ϵ from 0 to 0.1 for varying endowments. Selection strength is set to $s = 1,000$. Three points on the curve are highlighted: $\epsilon = 0.001$, $\epsilon = 0.019$, and $\epsilon = 0.1$. The point $\epsilon = 0.019$ is where the function attains its maximum, while the other two values are chosen arbitrarily for comparison. For these three values of ϵ , Fig. 4B shows the welfare achieved for all possible endowment distributions, still with $s = 1,000$. In this symmetric game, welfare is always maximized by $e = (0.5, 0.5)$. Finally, in Fig. 4C, ϵ is held constant at 0.019 while three different values of the selection strength s are shown for comparison. Fig. 4 D–F follow the same pattern, except that productivities are unequal and set to $r_1 = 1.3$, $r_2 = 1.9$. In Fig. 4D, the maximum is attained at $\epsilon = 0.043$. The same parameter values are used in [SI Appendix Figs. S1–S3](#).

Data, Materials, and Software Availability. Figures are based on simulation averages over many independent runs of the respective simulation. Results were analyzed and visualized with Python and Matlab R2023a. The computer code is published in (65). All other data are included in the manuscript and/or [SI Appendix](#).

ACKNOWLEDGMENTS. This work was supported by the European Research Council CoG 863818 (ForM-SMArt) (to K.C.) and the European Research Council Starting Grant 850529: E-DIRECT (to C.H.), the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie Grant Agreement #754411 and the French Agence Nationale de la Recherche (under the Investissement d'Avenir Programme, ANR-17-EURE-0010) (to M.K.).

Author affiliations: ^aInstitute of Science and Technology Austria, Klosterneuburg 3400, Austria; ^bSchool of Economics, The University of Queensland, Brisbane, QLD 4067, Australia; ^cMax Planck Research Group Dynamics of Social Behavior, Max Planck Institute for Evolutionary Biology, Plön 24306, Germany; and ^dInstitute for Advanced Study in Toulouse, Toulouse 31000, France

- B. Kerr, P. Godfrey-Smith, M. W. Feldman, What is altruism? *Trends Ecol. Evol.* **19**, 135–140 (2004).
- M. A. Nowak *et al.*, Evolving cooperation. *J. Theor. Biol.* **299**, 1–8 (2012).
- M. A. Nowak, Five rules for the evolution of cooperation. *Science* **314**, 1560–1563 (2006).
- R. Trivers, The evolution of reciprocal altruism. *Q. Rev. Biol.* **46**, 35–57 (1971).
- R. Axelrod, The emergence of cooperation among egoists. *Am. Polit. Sci. Rev.* **75**, 306–318 (1981).
- K. Sigmund, *The Calculus of Selfishness* (Princeton University Press, 2010), vol. 6.
- M. A. Nowak, K. Sigmund, Tit for tat in heterogeneous populations. *Nature* **355**, 250–253 (1992).
- M. R. Frean, The prisoner's dilemma without synchrony. *Proc. R. Soc. Lond. B* **257**, 75–79 (1994).
- T. Killingback, M. Doebeli, N. Knowlton, Variable investment, the continuous prisoner's dilemma, and the origin of cooperation. *Proc. R. Soc. Lond. B* **266**, 1723–1728 (1999).
- L. A. Imhof, M. A. Nowak, Stochastic evolutionary dynamics of direct reciprocity. *Proc. R. Soc. B: Biol. Sci.* **277**, 463–468 (2010).
- S. Kurokawa, J. Y. Wakano, Y. Ihara, Generous cooperators can outperform non-generous cooperators when replacing a population of defectors. *Theor. Popul. Biol.* **77**, 257–262 (2010).
- J. García, A. Traulsen, The structure of mutations and the evolution of cooperation. *PLoS ONE* **7**, e35287 (2012).
- W. H. Press, F. J. Dyson, Iterated Prisoner's Dilemma contains strategies that dominate any evolutionary opponent. *Proc. Natl. Acad. Sci. U.S.A.* **109**, 10409–10413 (2012).
- S. Van Segbroeck, J. M. Pacheco, T. Lenaerts, F. C. Santos, Emergence of fairness in repeated group interactions. *Phys. Rev. Lett.* **108**, 158104 (2012).
- A. Szolnoki, M. Perc, Defection and extortion as unexpected catalysts of unconditional cooperation in structured populations. *Sci. Rep.* **4**, 1–6 (2014).
- D. F. P. Toupo, D. G. Rand, S. H. Strogatz, Limit cycles sparked by mutation in the repeated prisoner's dilemma. *Int. J. Bifurcat. Chaos* **24**, 1430035 (2014).
- A. J. Stewart, J. B. Plotkin, Collapse of cooperation in evolving games. *Proc. Natl. Acad. Sci. U.S.A.* **111**, 17558–17563 (2014).
- A. J. Stewart, J. B. Plotkin, The evolvability of cooperation under local and non-local mutations. *Games* **6**, 231–250 (2015).
- Y. Dong, C. Li, Y. Tao, B. Zhang, Evolution of conformity in social dilemmas. *PLoS ONE* **10**, e0137435 (2015).
- L. Pan, D. Hao, Z. Rong, T. Zhou, Zero-determinant strategies in iterated public goods game. *Sci. Rep.* **5**, 1–10 (2015).
- S. K. Baek, H.-C. Jeong, C. Hilbe, M. A. Nowak, Comparing reactive and memory-one strategies of direct reciprocity. *Sci. Rep.* **6**, 1–13 (2016).
- A. McAvoy, C. Hauert, Autocratic strategies for iterated games with arbitrary action spaces. *Proc. Natl. Acad. Sci. U.S.A.* **113**, 3573–3578 (2016).
- J. G. Reiter, C. Hilbe, D. G. Rand, K. Chatterjee, M. A. Nowak, Crosstalk in concurrent repeated games impedes direct reciprocity and requires stronger levels of forgiveness. *Nat. Commun.* **9**, 555 (2018).
- G. Ichinose, N. Masuda, Zero-determinant strategies in finitely repeated games. *J. Theor. Biol.* **438**, 61–77 (2018).
- C. Hilbe, K. Chatterjee, M. A. Nowak, Partners and rivals in direct reciprocity. *Nat. Hum. Behav.* **2**, 469–477 (2018).
- L. Schmid, C. Hilbe, K. Chatterjee, M. A. Nowak, Direct reciprocity between individuals that use different strategy spaces. *PLoS Comput. Biol.* **18**, e1010149 (2022).
- O. P. Hauser, C. Hilbe, K. Chatterjee, M. A. Nowak, Social dilemmas among unequals. *Nature* **572**, 524–527 (2019).
- T. L. Cherry, S. Kroll, J. F. Shogren, The impact of endowment heterogeneity and origin on public good contributions: Evidence from the lab. *J. Econ. Behav. Organ.* **57**, 357–365 (2005).
- R. Merhej, F. P. Santos, F. S. Melo, F. C. Santos, "Cooperation between independent reinforcement learners under wealth inequality and collective risks" in *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems* (2021), pp. 898–906.
- M. Abou Chakra, S. Bumann, H. Schenk, A. Oschlies, A. Traulsen, Immediate action is the best strategy when facing uncertain climate change. *Nat. Commun.* **9**, 2566 (2018).
- T. Bergstrom, L. Blume, H. Varian, On the private provision of public goods. *J. Public Econ.* **29**, 25–49 (1986).

32. V. V. Vasconcelos, F. C. Santos, J. M. Pacheco, S. A. Levin, Climate policies under wealth inequality. *Proc. Natl. Acad. Sci. U.S.A.* **111**, 2212–2216 (2014).
33. Maria Abou Chakra and Arne Traulsen, Under high stakes and uncertainty the rich should lend the poor a helping hand. *J. Theor. Biol.* **341**, 123–130 (2014).
34. E. Akin, What you gotta know to play good in the iterated prisoner's dilemma. *Games* **6**, 175–190 (2015).
35. J. Zelmer, Linear public goods experiments: A meta-analysis. *Exp. Econ.* **6**, 299–310 (2003).
36. P. Dal Bó, G. R. Fréchet, On the determinants of cooperation in infinitely repeated games: A survey. *J. Econ. Lit.* **56**, 60–114 (2018).
37. E. Buckley, R. Croson, Income and wealth heterogeneity in the voluntary provision of linear public goods. *J. Public Econ.* **90**, 935–955 (2006).
38. S. P. Hargreaves, A. R. Heap, B. V. Stoddard, Endowment inequality in public goods games: A re-examination. *Econ. Lett.* **146**, 4–7 (2016).
39. S. Gächter, F. Mengel, E. Tsakas, A. Vostroknutov, Growth and inequality in public good provision. *J. Public Econ.* **150**, 1–13 (2017).
40. A. Nishi, H. Shirado, D. G. Rand, N. A. Christakis, Inequality and visibility of wealth in experimental social networks. *Nature* **526**, 426–429 (2015).
41. S. Gächter, Visible inequality breeds more inequality. *Nature* **526**, 333–334 (2015).
42. K. S. Chan, S. Mestelman, R. Moir, R. Andrew, Muller., Heterogeneity and the voluntary provision of public goods. *Exp. Econ.* **2**, 5–30 (1999).
43. C. Keser, A. Markstädter, M. Schmidt, Mandatory minimum contributions, heterogeneous endowments and voluntary public-good provision. *Games Econ. Behav.* **101**, 291–310 (2017).
44. J. Gross, S. Veistola, C. K. W. De Dreu, E. Van Dijk, Self-reliance crowds out group cooperation and increases wealth inequality. *Nat. Commun.* **11**, 5161 (2020).
45. A. Harrell, A. S. Greenleaf, Resource asymmetry reduces generosity and paying forward generosity, among the resource-advantaged and disadvantaged. *Soc. Sci. Res.* **109**, 102786 (2023).
46. J. C. Skewes, L. Nockur, National inequality, social capital, and public goods decision-making. *Curr. Res. Soc. Psychol.* **4**, 100112 (2023).
47. D. Melamed, B. Simpson, B. Montgomery, V. Patel, Inequality and cooperation in social networks. *Sci. Rep.* **12**, 6789 (2022).
48. A. J. Stewart, N. McCarty, J. J. Bryson, Polarization under rising inequality and economic decline. *Sci. Adv.* **6**, eabd4201 (2020).
49. L. Macchia, A. V. Whillans, The link between income, income inequality, and prosocial behavior around the world. *Soc. Psychol.* **52**, 375–386 (2022).
50. Yu. Fengyuan, J. Wang, J. He, Inequal dependence on members stabilizes cooperation in spatial public goods game. *Chaos Solit. Fractals* **165**, 112755 (2022).
51. L. Nockur, L. Arndt, J. Keller, S. Pfattheicher, Collective choice fosters sustainable resource management in the presence of asymmetric opportunities. *Sci. Rep.* **10**, 10724 (2020).
52. W. Barfuss, J. F. Donges, V. V. Vasconcelos, J. Kurths, S. A. Levin, Caring for the future can turn tragedy into comedy for long-term collective action under risk of collapse. *Proc. Natl. Acad. Sci. U.S.A.* **117**, 12915–12922 (2020).
53. R. Selten, *Reexamination of the Perfectness Concept for Equilibrium Points in Extensive games* (Springer, 1988).
54. D. Fudenberg, J. Tirole, *Game Theory* (MIT Press, Cambridge, MA, 1991), vol. 86.
55. M. C. Couto, S. Giarmo, C. Hilbe, Introspection dynamics: A simple model of counterfactual learning in asymmetric games. *New J. Phys.* **24**, 063010 (2022).
56. H. Zhang, Errors can increase cooperation in finite populations. *Games Econ. Behav.* **107**, 203–219 (2018).
57. K. Hoff, A. B. Lyon, Non-leaky buckets: Optimal redistributive taxation and agency costs. *J. Public Econ.* **58**, 365–390 (1995).
58. G. C. Loury, Intergenerational transfers and the distribution of earnings. *Econometrica* **49**, 843–867 (1981).
59. A. V. Banerjee, A. F. Newman, Occupational choice and the process of development. *J. Polit. Econ.* **101**, 274–298 (1993).
60. R. Benabou, Unequal societies: Income distribution and the social contract. *Am. Econ. Rev.* **91**, 96–129 (2000).
61. O. Galor, J. Zeira, Income distribution and macroeconomics. *Rev. Econ. Stud.* **60**, 35–52 (1993).
62. T. Piketty, The dynamics of the wealth distribution and the interest rate with credit rationing. *Rev. Econ. Stud.* **64**, 173–189 (1997).
63. R. J. Dwyer, E. W. Dunn, Wealth redistribution promotes happiness. *Proc. Natl. Acad. Sci. U.S.A.* **119**, e2211123119 (2022).
64. C. S. Holling, Resilience and stability of ecological systems. *Annu. Rev. Ecol. Syst.* **4**, 1–23 (1973).
65. V. Hübner, M. Kleshnina, Computer code for "Efficiency and resilience of cooperation in asymmetric social dilemmas." Zenodo. <https://doi.org/10.5281/zenodo.10639168>. Deposited 9 February 2024.