PNAS Nexus, 2025, **4**, pgaf154
https://doi.org/10.1093/pnasnexus/pgaf154
Advance access publication 10 May 2025
Research Report

Stable strategies of direct and indirect reciprocity across all social dilemmas

Valentin Hübner[a], Laura Schmid[b], Christian Hilbe[c,d,1] and Krishnendu Chatterjee[a,*,1]

[a]Institute of Science and Technology Austria, Klosterneuburg 3400, Austria
[b]Kim Jaechul Graduate School of AI, Korea Advanced Institute of Science and Technology, Seoul 02455, Republic of Korea
[c]Max Planck Research Group on the Dynamics of Social Behaviour, Max Planck Institute for Evolutionary Biology, Plön 24306, Germany
[d]Interdisciplinary Transformation University IT:U, Linz 4020, Austria
*To whom correspondence should be addressed: Email: krishnendu.chatterjee@ista.ac.at
[1]C.H. and K.C. contributed equally to this work
Edited By David G. Rand

Abstract

Social dilemmas are collective-action problems where individual interests are at odds with group interests. Such dilemmas occur frequently at all scales of human interactions. When dealing with collective-action problems, people often act reciprocally. They adjust their behavior to match the previous behavior of the recipient. The literature distinguishes two kinds of reciprocity. According to direct reciprocity, individuals react to their immediate experiences with the recipient. They are more likely to cooperate if the recipient previously cooperated with them. According to indirect reciprocity, individuals react to the recipient's general behavior, irrespectively of whether or not they benefited directly. In practice, the two kinds of reciprocity are often intertwined; people typically base their decisions on both direct experiences and indirect observations. Yet only recently have researchers begun to explore how the two kinds of reciprocity interact. So far, this research only addresses a single type of social dilemma, the donation game, where the effects of individual behaviors are independent. Instead, here we allow for all pairwise social dilemmas. By applying novel techniques to generalize the theory of zero-determinant strategies, we establish an important proof of principle: In all social dilemmas, socially optimal outcomes can be sustained as an equilibrium, using either direct or indirect reciprocity, or arbitrary mixtures thereof. These results neither require games to be repeated infinitely often, nor that individual opinions are synchronized. In this way, we considerably generalize the scope of models of reciprocity, and we build further bridges between the literatures on direct and indirect reciprocity.

Keywords: social dilemmas, direct reciprocity, indirect reciprocity, equalizer strategies

Significance Statement

Social dilemmas are decision-making problems in which there is a conflict between collective and individual interests. Two prominent approaches to resolve such dilemmas are direct and indirect reciprocity. In direct reciprocity, individuals react to their personal experiences with others. In indirect reciprocity, they act based on others' reputations. Although the two forms of reciprocity are intertwined in practice, most models study them in isolation. Our work combines both forms in a framework that applies to all social dilemmas. Based on this framework, we provide a general existence proof. We show that full cooperation can always be sustained as a Nash equilibrium, independent of whether games are discounted, whether opinions are synchronized, and whether individuals use direct or indirect reciprocity.

Introduction

People frequently encounter situations in which individually optimal behaviors diminish the welfare of others. Such social dilemmas may, for example, lead individuals to put too little effort into group projects, or to overuse public resources (1–4). These types of conflict can be analyzed using the mathematical framework of (evolutionary) game theory (5–7). This framework provides tools to describe individuals who, consciously or subconsciously, make decisions that affect others' well-being. In particular, this literature describes several mechanisms that help individuals to cope with their social dilemmas (8). One prominent mechanism, especially in the context of pairwise interactions, is reciprocity. According to this mechanism, individuals have more of an incentive to act in the interests of others if their actions today may be reciprocated in the future.

The literature on evolutionary game theory distinguishes several types of reciprocity. One type is direct reciprocity (9–12). Here, individuals decide how to act based on their previous

Competing Interest: The authors declare no competing interests.
Received: March 31, 2025. **Accepted:** April 8, 2025

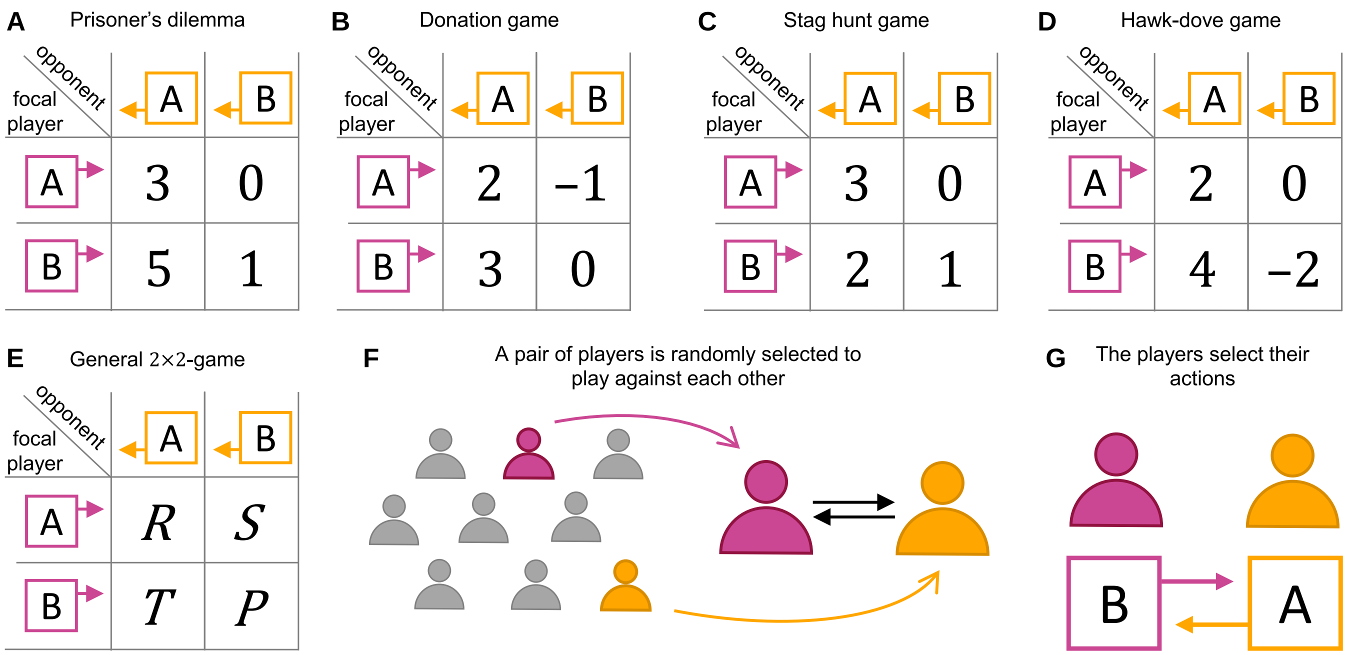

Fig. 1. Illustration of the game dynamics. We consider populations of players who engage in pairwise social dilemmas with two actions *A* and *B*. A) In the prisoner's dilemma, the best outcome for the group is for both players to choose *A*. Yet individually, *B* is a dominant action. B) The donation game is a special case of a prisoner's dilemma. Here, action *A* can be interpreted as paying a cost $c > 0$ for the coplayer to receive a benefit $b > c$. Action *B* means to do nothing. In the depicted example, $b = 3$ and $c = 1$. C) The stag hunt game captures a dilemma in which players may fail to coordinate on the most profitable equilibrium. D) In the hawk-dove game, there are two (pure) equilibria. In each equilibrium, one player chooses *A* and the other one *B*. Each player prefers to be the one who chooses *B*. E) In general, payoffs are often denoted by the letters *R*, *S*, *T*, and *P*, respectively. F) The game unfolds over many rounds. Each round, two players are randomly drawn to interact with each other in the given social dilemma. G) Each of the players chooses action *A* or action *B*. Their choice might depend on their coplayer's previous interactions.

experience with the respective interaction partner. That is, when Alice decides how to treat Bob, she considers how Bob treated her in the past. Such conditional behaviors have primarily been explored in the context of the prisoner's dilemma (13–16) (Fig. 1A). In this game, the socially optimal choice of cooperation is dominated by defection. However, once individuals interact repeatedly, reciprocal strategies such as Tit-for-Tat (9) can help sustain cooperation. Even though the repeated prisoner's dilemma has been the main model to study direct reciprocity, the very same mechanism can also be effective in weaker forms of social conflict (17, 18).

Another type of reciprocity is indirect reciprocity (19–22). Here, when Alice decides how to treat Bob, she takes into account Bob's overall behavior, including how he acted towards Charlie or Dave. That is, she takes into account Bob's general reputation. Unlike direct reciprocity, this mechanism does not require repeated interactions among the same two individuals. It merely requires that individuals repeatedly interact within a larger community. With a few exceptions (23, 24), researchers study this type of reciprocity with an even more restricted type of social dilemma, the donation game (a special case of the prisoner's dilemma, see Fig. 1B). This literature suggests that cooperation can be sustained with a variety of strategies, most notably the "leading-eight" norms (25, 26).

Although direct and indirect reciprocity are based on a similar premise, most theoretical and experimental studies (27–30) either consider one or the other. Such a reductionist approach is useful to clarify whether either mechanism can be effective on its own. At the same time, however, it renders many interesting research questions infeasible. For example, such models cannot explain how individuals would cope with conflicting pieces of evidence (e.g. when Alice's personal impression of Bob runs counter to his public reputation (31, 32)). Similarly, such models cannot explain why in direct reciprocity, cooperation can be maintained with comparably simple strategies, whereas indirect reciprocity seems to require strategies of greater complexity (33–38). Only more recently, researchers have begun to describe different types of reciprocity within a single framework (39–42). The corresponding studies explore when people would rather adopt one type of reciprocity instead of the other. Unfortunately, however, these studies are restricted to the analysis of donation games only. As a result, they cannot capture synergistic interactions, as in the stag hunt game (43, 44) (Fig. 1C). Similarly, they cannot capture cases in which one individual's cooperation crowds out the need for others to cooperate, as in the volunteer's dilemma (45), the snowdrift game (46) or other classes of hawk-dove games (Fig. 1D). To describe the effects of reciprocity in full generality, it takes models that allow for all kinds of social dilemmas. We present such a model herein.

Such a generalization is not straightforward. In donation games, the payoff consequences of one individual's cooperation are independent of whether or not their interaction partner cooperates too. This independence allows researchers to compute the players' payoffs explicitly, by solving a low-dimensional system of linear recursions (42). Beyond donation games, this simple recursion no longer applies. Hence, an analysis of direct and indirect reciprocity across all social dilemmas requires a different set of proof techniques, which we summarize below (and which we discuss in full detail in the Supplementary Information).

To characterize whether socially optimal outcomes can be sustained with either direct or indirect reciprocity (or both), we extend the notion of so-called *equalizer* strategies. These strategies have been first introduced in the context of direct reciprocity (47, 48). By implementing an equalizer strategy, individuals can unilaterally set their opponent's payoff to a fixed value. That is, opponents always get the same payoff, irrespective of their own behavior. Once all other players adopt an equalizer strategy, the remaining player thus neither has an advantage, nor a disadvantage, from deviating. This property makes equalizers a useful tool to prove the abstract existence of Nash equilibria. The resulting set of equalizers includes several well-known strategies, such as Generous Tit-for-Tat (49, 50) and generalizations thereof (51). By building on these ideas, Schmid et al. (42) have shown that equalizers can also be used to sustain full cooperation in models of indirect reciprocity—provided the game at hand is a donation game. Herein, we characterize when such equalizers exist in arbitrary pairwise social dilemmas, for both direct and indirect reciprocity (and arbitrary mixtures). Along the way, we also prove that in general, equalizers need to be more complex than previously appreciated (in technical terms: in donation games, equalizers can be implemented with simple *reactive* strategies (42). For more general social dilemmas, it takes the richer set of *memory-1* strategies instead).

Our model combines direct and indirect reciprocity within a single framework, irrespective of the considered social dilemma. This framework can serve as an important bridge to transfer insights from one field to another. Herein, we use this bridge, for example, to incorporate the well-known memory-one strategies from direct reciprocity into models of indirect reciprocity. In this way, we can prove that even in indirect reciprocity, fully cooperative outcomes can be sustained as a Nash equilibrium. In the past, such rigorous results for indirect reciprocity have been difficult to establish; these difficulties have been especially pronounced in the case of "private information," when individuals are allowed to disagree on each others' reputations (52). Instead, here we prove the existence of such Nash equilibria for arbitrary social dilemmas, and for players who are allowed to discount the future—even when information is private.

Results

Game setup

We consider a population of n individuals, referred to as players. These players repeatedly interact in a pairwise interaction. More specifically, in each round, two players are selected at random (Fig. 1F). Each player then chooses one of two actions, A or B (Fig. 1G). Their choices determine the payoffs they get, according to the given payoff matrix (Fig. 1A–E). All other population members observe the interaction, but they may independently misperceive each player's action with some probability $\varepsilon < \frac{1}{2}$. Instead of correctly identifying a player's action as, say, action A, they perceive it to be B, or vice versa. After the interaction has taken place, there is another round with continuation probability d. In that case, a new pair of players is randomly drawn to interact with one another. With the converse probability $1 - d$, the game ends. The total payoff of each player is defined as the sum of the payoffs they obtained in each round, times a normalization constant. Equivalently, one may also interpret this setup as an infinitely repeated interaction in which players discount future rounds by a constant factor d. In the limit $d \to 1$, we recover the classical case of an infinitely repeated game without discounting, as for example in Press and Dyson (48); see Methods for details.

The exact nature of the game played each round depends on the four entries R, S, T, P of the payoff matrix (Fig. 1E). In the following, we are particularly interested in games that can be interpreted as social dilemmas. Based on the "individual-based" interpretation in Kerr et al. (53), this means payoffs satisfy the following constraints. First, players prefer mutually choosing A to mutually choosing B, such that $R > P$ (except for the degenerate case of $R = P$, this assumption is without loss of generality; otherwise we just need to relabel the two actions). Second, players always prefer their coplayer to choose A, implying $R > S$ and $T > P$. Third, in mixed pairs, the player who chooses B gets the higher payoff, such that $T > S$. Together these assumptions ensure that on a collective level, individuals have some incentive to choose A; yet on an individual level, they may want to choose B. Accordingly, we interpret action A as cooperation, and we associate B with defection (however, we use the more neutral letters A and B, rather than the usual letters C and D, to highlight that our framework is not restricted to the prisoner's dilemma).

The notion of a social dilemma captures several classical games. (i) In the prisoner's dilemma, payoffs satisfy the inequalities $T > R > P > S$ and $2R > T + S$, as in Fig. 1A. (ii) The donation game additionally requires $R + P = S + T$, see Fig. 1B. Such games are sometimes called "additive" (54, 55). (iii) The stag hunt game satisfies $R > T > P > S$, as depicted in Fig. 1C. (iv) Finally, the hawk-dove game satisfies $T > R > S > P$, as in Fig. 1D. The exact payoff ranking determines the severity of the dilemma. Among the above examples, players arguably face the strongest conflict between cooperation and defection in the prisoner's dilemma and the donation game. However, also the other two games entail some conflict. Players may either have difficulties to coordinate on the equilibrium that is better for both (as in the stag hunt game), or they may prefer different equilibria altogether (as in the hawk-dove game).

Reactive and memory-1 strategies

When playing the above games, players make their decisions based on their strategies. Strategies are recipes that tell the player what to do, depending on the outcome of previous interactions.

In order to allow for an explicit analysis, researchers often consider a restricted space of strategies. For example, Schmid et al. (42) consider strategies of the form $\sigma = (p_0, p_A, p_B, \lambda)$. Here, the first parameter p_0 is a player's cooperation probability against an unknown coplayer. The next two parameters p_A (p_B) give the player's cooperation probability against a coplayer who cooperated (defected) in their last relevant interaction. Finally, the parameter λ determines which previous interactions of the coplayer are deemed relevant. When $\lambda = 0$, only direct interactions matter. For example, if Bob previously defected against Alice (played B), but then cooperated with Charlie (played A), Alice would use cooperation probability p_B against Bob. That is, Alice implements a strategy of direct reciprocity (Fig. 2A). In contrast, when $\lambda = 1$, players take into account all their coplayers' interactions equally, even interactions with third parties. As a result, such players base their decision on the very last action of the coplayer, independently of whether or not they were personally involved. In the above example, if Bob defected against Alice (played B), but then cooperated with Charlie (played A), Alice's cooperation probability against Bob is p_A. Now, Alice uses a strategy of indirect reciprocity (Fig. 2B). The model also allows for intermediate values of $\lambda \in (0, 1)$. In that case, Alice takes into account third-party interactions with probability λ (Fig. 2C). As the above strategies merely respond to the *coplayer's* previous behavior, they are called *reactive* (7) (Fig. 2D). The set of reactive strategies includes Always Cooperate $\sigma = (1, 1, 1, \lambda)$, Tit-for-Tat $\sigma = (1, 1, 0, 0)$, and its indirect reciprocity analog Simple Scoring (56) $\sigma = (1, 1, 0, 1)$, among others.

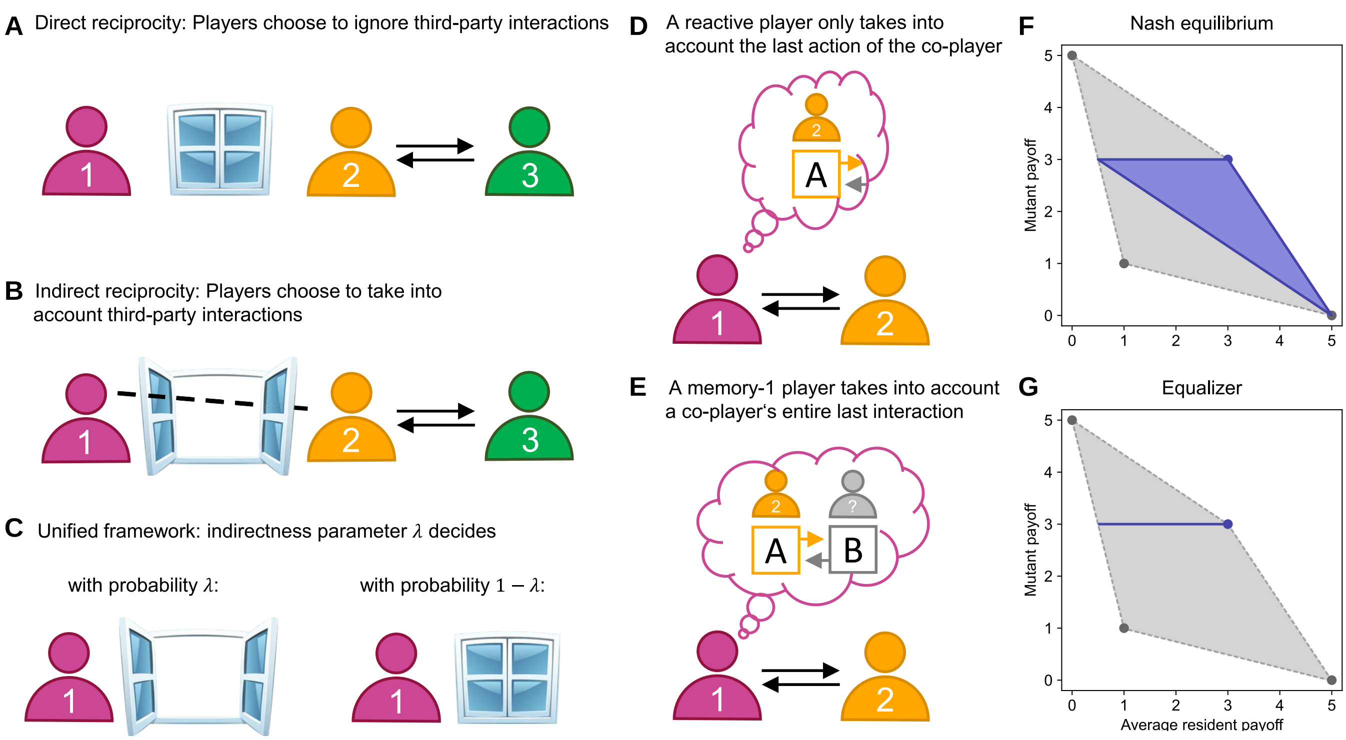

Fig. 2. Strategies of direct and indirect reciprocity. Strategies of reciprocity differ in whether or not players (here, Player 1) take into account third-party interactions (here, between players 2 and 3). A) According to direct reciprocity, Player 1 ignores third-party interactions. B) According to indirect reciprocity, Player 1, takes such third-party interactions into account. C) Our framework also allows for intermediate cases, where player 1 considers third-party interactions with some fixed probability λ. D) For our model, we consider strategies of different complexity. When using a reactive strategy, players condition their behavior on the last observed action of the coplayer. E) When using a memory-1 strategy, they condition their behavior on the entire outcome of the coplayer's last interaction. This also involves the action of the coplayer's last opponent. F) A strategy is a Nash equilibrium if it is a best response to itself. The effect of such strategies can be represented graphically. The large area shaded in gray represents all feasible payoffs in the respective game. The smaller area, shown in blue, represents the payoffs that are still feasible if every player adopts the same fixed resident strategy, except for one deviating mutant. In this example, the resident strategy yields a payoff of three against itself (indicated by the upper right dot). According to the blue area, no mutant strategy yields a higher payoff. Hence, the given resident strategy is a Nash equilibrium. G) An equalizer strategy is a special case of a Nash equilibrium. Here, the mutant's payoff is always the same, regardless of the mutant's strategy.

In the context of direct reciprocity, it is also common to consider a slightly more general strategy set, called *memory-1 strategies*. Here, a player does not only take into account the coplayer's last action. Rather the player takes into account the entire context of the coplayer's previous interaction, including the action of the co-player's opponent (Fig. 2E). Memory-1 strategies take the form $\sigma = (p_0, p_{AA}, p_{AB}, p_{BA}, p_{BB}, \lambda)$. The interpretation of the entries p_0 and λ is the same as before. However, now p_{xy} is a player's cooperation probability given that in the coplayer's last relevant interaction, the coplayer used action y whereas the coplayer's opponent used action x. Within the set of memory-1 strategies we can represent reactive strategies as those strategies for which $p_{AA} = p_{BA}$ and $p_{AB} = p_{BB}$. Here, only the coplayer's last relevant action matters. A well-known example of a nonreactive memory-1 strategy is Win-Stay Lose-Shift (57), $\sigma = (1, 1, 0, 0, 1, 0)$. Here, a player would only cooperate with a coplayer if in their previous joint interaction either both cooperated, or no one did (57).

Partner strategies

In the following, we are interested in whether mutual cooperation can be sustained by either direct or indirect reciprocity. To this end, we study strategies σ with two properties. First, the strategy ought to be *nice* (9). That is, if σ is adopted by everyone, the entire population cooperates indefinitely in the absence of errors. Second, the respective strategy ought to be a Nash equilibrium: if adopted by everyone, no single player has an incentive to deviate (Fig. 2F). In the context of direct reciprocity, strategies that satisfy both properties are called *partners* (58). The answer to the question whether partners exist turns out to be trivial in the stag hunt game or in the so-called harmony game (17). In those games, payoffs satisfy $R > T$. Therefore, mutual cooperation is a Nash equilibrium even if the game is only played once. It trivially follows that mutual cooperation can also be sustained within our repeated setup—players merely need to use the strategy Always Cooperate. In the following, we will thus focus more on the other two game classes, the prisoner's dilemma and the hawk-dove game.

To show existence of partner strategies in those games, we characterize a particular subset of Nash equilibria, those based on equalizer strategies. Such strategies do not only ensure that no player can unilaterally improve their payoff; they ensure every deviating player gets the same payoff (Fig. 2G). For direct reciprocity ($\lambda = 0$), the existence of equalizers has been shown by Boerljist et al. (47) and Press and Dyson (48). Their result applies to the infinitely repeated prisoner's dilemma without errors ($d = 1$ and $\varepsilon = 0$). For indirect reciprocity ($\lambda = 1$), the existence of equalizers follows from the work of Schmid et al. (42), but only for the restrictive case of donation games (but arbitrary d and ε). Instead, here we characterize equalizers for all social dilemmas, for direct and indirect reciprocity, and for all continuation probabilities and error rates. All details and proofs are in the Supplementary Information. Below we summarize the respective results. For a visual representation of previous work and our contribution, see Fig. 3.

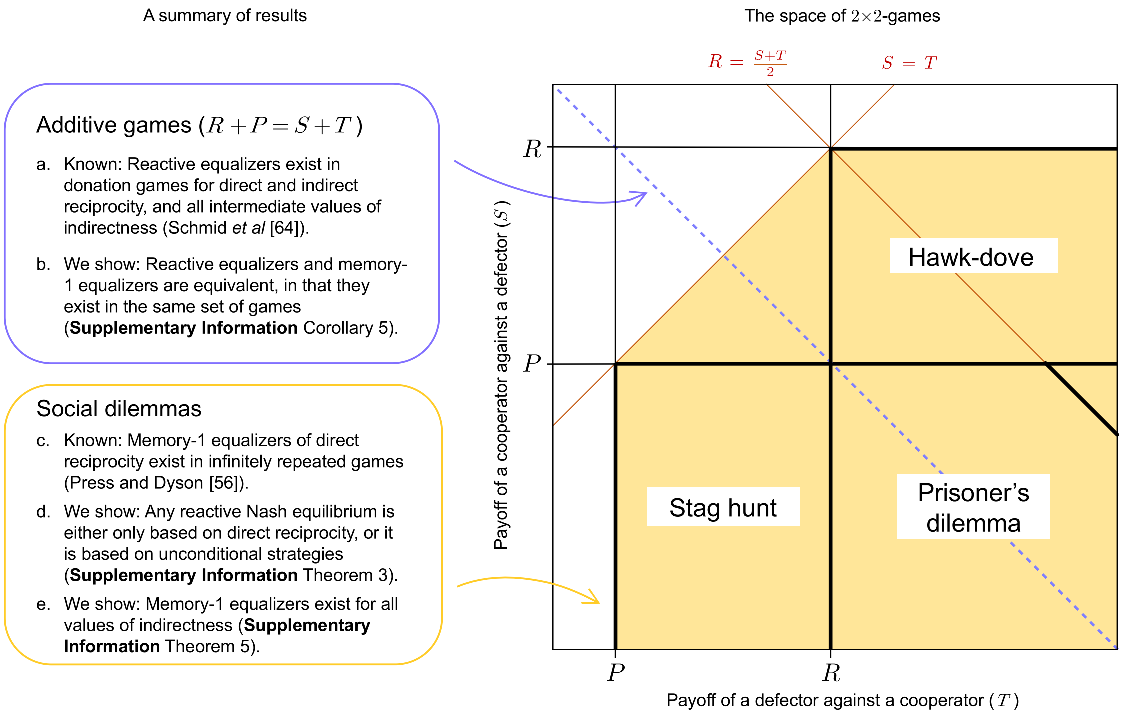

Fig. 3. Summary of our results. Herein, we derive results on the existence of "equalizer strategies" in pairwise social dilemmas. The right half of the figure represents the space of all such dilemmas graphically. Here, we keep the two payoff parameters R and P fixed (with $R > P$). We vary the remaining payoffs T and S. The region shaded in orange indicates all games that satisfy the conditions of a social dilemma. The blue dashed line indicates the subspace of "additive games," which includes the donation game. All previous models that combine direct and indirect reciprocity focus on this blue subspace (39–42). For the general space of social dilemmas, the existence of equalizers has only been established for direct reciprocity (48), but not for indirect reciprocity or any mixtures.

Reactive strategies in the donation game

To better motivate our contribution, let us first recapitulate the results of Schmid et al. (42). They considered a similar setup as ours, but restricted to the donation game and to players with reactive strategies. Payoffs of the donation game are given by $R = b - c$, $S = -c$, $T = b$, and $P = 0$, where b and c are the benefit and cost of cooperation. Schmid et al. show that for full cooperation to be sustainable, the pairwise continuation probability δ needs to be sufficiently large (this is the probability that a given pair of players will interact again given it just interacted; this probability is directly related to the population-wide probability d, see Supplementary Information). The exact threshold for δ depends on the players' indirectness parameter λ. For direct ($\lambda = 0$) and indirect reciprocity ($\lambda = 1$), the respective thresholds are

$$\delta_0 = \frac{c}{b} \quad \text{and} \quad \delta_1 = \frac{c}{b + (n-2)((1-2\varepsilon)b - c)}. \qquad (1)$$

In particular, indirect reciprocity makes it easier to sustain cooperation (compared to direct reciprocity) if the population is large and errors are rare. Either way, once the respective threshold is reached, mutual cooperation can be enforced with an equalizer strategy. In case of direct reciprocity, the respective equalizer is Generous Tit-for-Tat (50). In case of indirect reciprocity, it is Generous Scoring (42).

To derive the above results, both the restriction to donation games and to reactive strategies turns out to be crucial. Because the donation game satisfies the additivity property $R + P = S + T$, the payoff of each player can be decomposed into a sum of two terms. The first term only depends on the statistical distribution of the coplayers' actions (affecting whether or not I receive a benefit b). The other term only depends on the statistical distribution of the own action (affecting whether or not I pay the cost c). That is, the players' actions affect payoffs independently. Furthermore, for reactive strategies, the distribution of a player's own actions affects the distribution of the co-players' actions by a linear relationship. Based on these two observations, one can derive a simple linear recursion for the players' likelihood to cooperate with each other in any given round. With this recursion, it becomes straightforward to compute payoffs. Unfortunately, once either the game is nonadditive, or players use more complex strategies, the above approach is no longer viable. Hence, for the results below we rely on proof techniques that do not require us to compute the players' payoffs explicitly.

Memory-1 strategies in the donation game

To make progress, we first explore whether cooperation in the donation game is easier to sustain when players are allowed to use memory-1 strategies. More specifically, we ask whether there are memory-1 equalizers that can sustain full cooperation even when the respective condition in (1) is violated. The answer is negative. We find that for all game parameters and any indirectness λ, memory-1 equalizers exist if and only if reactive equalizers exist (Supplementary Information Corollary 5). This result resonates with earlier work on direct reciprocity. For the infinitely repeated donation game, it was shown that reactive strategies can enforce all linear payoff relationships that are theoretically possible (59). Thus at least in the donation game, allowing for more complex strategies does not provide any additional advantage with respect to implementing equalizer strategies (Fig. 3B).

Reactive strategies in general social dilemmas

Given the strong properties of reactive strategies in donation games, we ask whether they can also sustain full cooperation in other social dilemmas. Surprisingly, the answer is negative. To describe this result more formally, we introduce the notion of a degenerate strategy. A reactive strategy is degenerate if it exclusively relies on direct reciprocity ($\lambda = 0$), or if it acts unconditionally ($p_A = p_B$). That is, degenerate strategies completely ignore any third-party interactions. Similarly, we say an equilibrium is degenerate if it requires players to use degenerate strategies. Using this notion, we can formulate the main result of this section as follows: In any nonadditive game with more than two players ($n > 2$) and positive error rates ($\varepsilon > 0$), any Nash equilibrium in reactive strategies is degenerate. In other words, if players are in a Nash equilibrium that entails at least some indirect reciprocity ($\lambda > 0$), players must be using unconditional strategies such as Always Defect. (The above result also implies that for $\lambda > 0$, there are usually no equalizer strategies, because unconditional strategies are in general not equalizers (59).) This finding suggests that earlier results on the donation game (42) are sensitive to the exact payoff values. Once the payoff matrix is slightly perturbed, reactive partner strategies that entail some indirect reciprocity cease to exist (Fig. 3D).

The proof of the above result is constructive: We show that for any such reactive resident strategy, one can construct a deviating strategy that gets a strictly higher payoff. Interestingly, the deviating strategy is not reactive. Rather, it is a higher-memory strategy that takes into account the joint distribution of previous actions across different pairs of players. We show that such strategies have a payoff advantage when the process involves at least some randomness (e.g. when there are errors). We provide a description of the deviation strategies in the Methods, and a proof of their superiority in the Supplementary Information.

Memory-1 strategies in general social dilemmas

The above result raises the question whether beyond the simple donation game, nondegenerate equalizers exist at all. To explore that question, we search the space of memory-1 strategies. There, we find that the answer is positive. For any social dilemma and any indirectness λ, there exist equalizer strategies for sufficiently large continuation probabilities and sufficiently small error rates (Fig. 3E, see Supplementary Information). Similar to (1) for the donation game, the minimum continuation probability can be computed explicitly. For example, assuming $\varepsilon = 0$ and $T > R$, we find that fully cooperative equalizers with indirectness λ exist if and only if the pairwise continuation probability exceeds the threshold [2],

$$\delta_\lambda = \left(1 + (1 + (n-2)\lambda)\,\frac{\min\{T, R\} - \max\{P, S\}}{\max\{|T - R|, |P - S|\}}\right)^{-1}. \qquad (2)$$

A few remarks are in order. First, for social dilemmas, this threshold is strictly smaller than one. Hence, the condition can be satisfied for sufficiently large δ. Second, for any population size $n > 2$, it is easy to verify that threshold (2) is strictly lower for indirect reciprocity than for direct reciprocity. This is a consequence of our assumption that there are no errors. Once the error rate becomes positive, direct reciprocity may become the more favorable mechanism for full cooperation (Fig. 4). Third, in the special case of the donation game, the threshold simplifies to the following values for direct ($\lambda = 0$) and indirect reciprocity ($\lambda = 1$):

$$\delta_0 = \frac{c}{b} \quad \text{and} \quad \delta_1 = \frac{c}{b + (n-2)(b-c)}. \qquad (3)$$

That is, we recover the earlier conditions in (1) for reactive strategies for $\varepsilon = 0$.

In the more general case of an arbitrary prisoner's dilemma and of the hawk-dove game (with $T > R$), we show that once the condition (2) is satisfied, one can always find equalizer strategies that enforce the mutual cooperation payoff R (see Supplementary Information, Proposition 5). Again, our proof is constructive. In the Methods, we provide an algorithm that produces an optimal equalizer strategy for all social dilemmas (even for positive error rates). For games with $T > R$, this algorithm produces nice strategies (i.e. $p_0 = p_{AA} = 1$). Together with our earlier observation that cooperation is trivial to sustain in the stag hunt and the harmony game (with $T < R$), we conclude that stable cooperation can always be achieved with memory-1 strategies, based on direct or indirect reciprocity, or any arbitrary mixture of the two.

Overall, the above results represent a considerable generalization of previous work. We recover the seminal results of Press & Dyson (48), when we restrict our framework to direct reciprocity in infinitely repeated games ($\lambda = 0$ and $d = 1$). Similarly, we recover the results of Schmid et al. (42), when we restrict our framework to reactive strategies, and to donation games only (see Methods for details).

Simulation results

To further illustrate the above results, we have explored the game dynamics when $n - 1$ players act according to a given equalizer strategy (produced by Algorithm 1 in the Methods section). For the remaining player, we have sampled $N = 100$ random "mutant" strategies. To approximate the players' resulting payoffs, we simulated many independent instances of the game dynamics, separately for each mutant strategy (in contrast to previous work on reactive strategies in donation games (42), there is no known formula to compute the players' payoffs explicitly). Figure 5 shows the results. We depict the residents' and the mutant's average payoff for three different social dilemmas (the prisoner's dilemma, the stag hunt game, and the hawk-dove game). In each case, we find that the simulated payoffs indeed form a straight horizontal line, the characteristic property of an equalizer strategy (Fig. 2G). In particular, in each case the produced resident strategy is a Nash equilibrium: Once adopted by everyone, no mutant strategy has a selective advantage.

Discussion

Direct and indirect reciprocity are important determinants of human behavior in social dilemmas (8). They are arguably among the key mechanisms to explain our exceptionally high cooperation rates (60). Yet despite the many similarities between the two kinds of reciprocity, they are typically studied independently. Even worse, respective models often differ substantially. Models of direct reciprocity tend to study the prisoner's dilemma (9–12), whereas indirect reciprocity models are often based on the narrower class of donation games (19–22). Similarly, individuals in direct reciprocity models are typically assumed to adopt reactive or memory-1 strategies (7, 61). In contrast, models of indirect reciprocity focus on subsets of "third-order social norms" (26), which do not map easily onto either class of direct reciprocity strategies. These differences make it difficult to compare the two mechanisms directly. Moreover, they make it difficult to generalize insights from one field to the other. To address these problems, we join recent efforts to study a unified framework (39–42), in which individuals themselves choose which kind of reciprocity they use.

Previous models that combine direct and indirect reciprocity are based on the smallest common denominator of the two

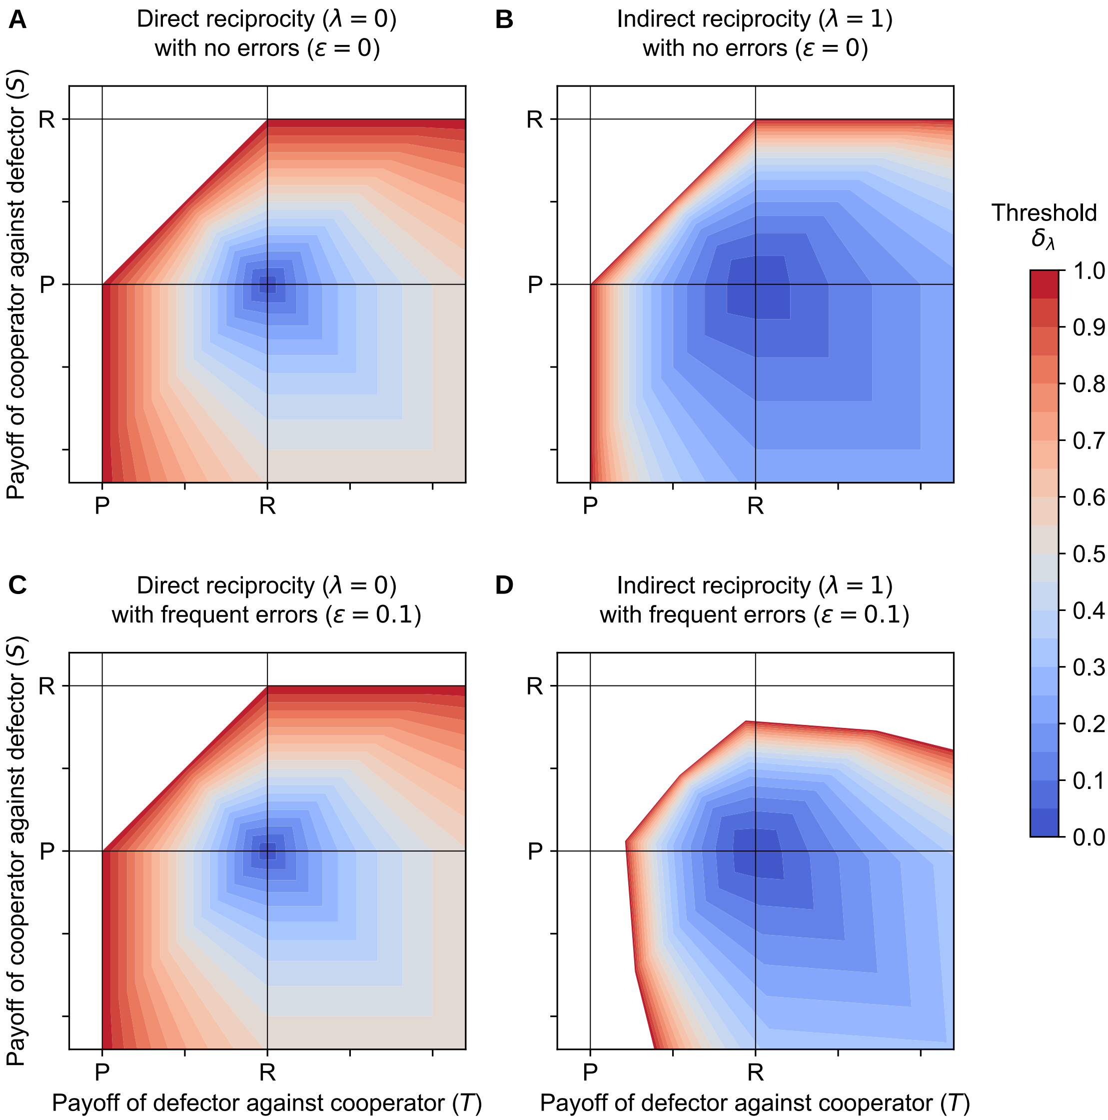

Fig. 4. Feasibility of equalizer strategies across all social dilemmas. We graphically represent whether or not equalizer strategies exist. To this end, we consider four different cases. The cases depend on whether individuals use direct (left) or indirect reciprocity (right), and on whether or not there are errors (top vs. bottom). In each case, we consider the space of all social dilemmas (as in Fig. 3). For each possible game, we depict how large the pairwise continuation probability δ needs to be for equalizers to exist. Low values of δ (blue, near the center) indicate that the conditions for equalizers are easy to satisfy. Higher values of δ (red, near the boundary) suggest that equalizer strategies only exist for rather high continuation probabilities. A, B) The figure shows that without errors, indirect reciprocity is more favorable to the existence of equalizers. C, D) Once third-party observations are subject to errors, there are regions in which direct reciprocity allows for equalizers whereas indirect reciprocity does not. An expanded view of the same graph is shown in Fig. S1 in the Supplementary Information.

literatures, the donation game. This game is the simplest metaphor of cooperation: individuals pay some cost to provide a benefit to someone else. This simplicity promotes a mathematical analysis, and it permits an intuitive interpretation of the results. However, this game rules out possible interdependencies between the individuals' actions. Neither is it particularly beneficial if individuals cooperate at the same time, nor is it particularly damaging if they all defect simultaneously. This assumption makes it impossible to study games in which mutual cooperation yields synergistic benefits, as in the stag hunt game. Similarly, it rules out interactions in which individual actions are strategic substitutes, as the volunteer's dilemma (45) or the snowdrift game (46). Even among all prisoner's dilemmas, donation games only represent a negligible subset of measure zero. These considerations highlight a need to study models that allow for more general types of social dilemmas. We present such a model herein.

We use this model to characterize strategies that can sustain full cooperation. Our results show that these strategies do not need to be overly complex. Instead, it suffices that individuals take into account the last interaction of the respective group member (e.g. to consider memory-1 strategies). Our results also show that simpler strategies (reactive strategies) in general do not suffice to support cooperation. In fact, the only domain in which these strategies suffice are the donation games considered earlier (39–42). Together, these two observations represent a nice characterization of the complexity of strategies that is necessary and sufficient to ensure stable cooperation for all social dilemmas.

While the basic setup we consider is thus similar to earlier unified frameworks of reciprocity (39, 41, 42), the mathematical tools we apply are vastly different. Earlier work exploited the advantage that payoffs of reactive players in the donation game can be

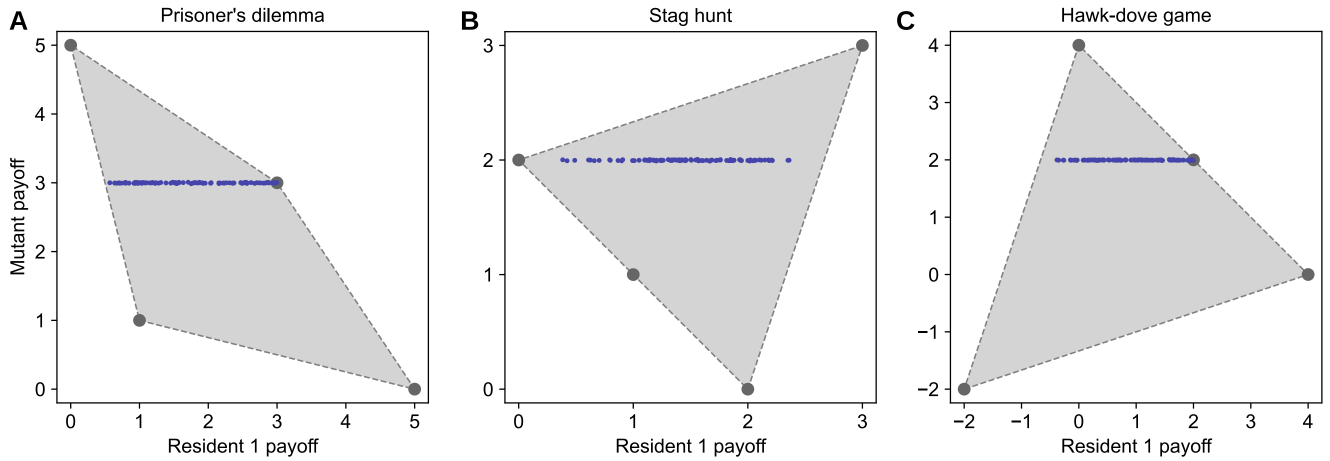

Fig. 5. Simulation of equalizer strategies. We consider three social dilemmas in which $n-1$ residents use a fixed equalizer strategy. For the remaining player, we randomly sample $N = 100$ mutant strategies. For each mutant strategy, we simulate the resulting game dynamics. Based on these simulations, we compute the average payoff in pairwise interactions between the deviating mutant player and a given resident player ("Resident 1," see Methods for details). These pairwise payoffs are depicted as small blue dots. As expected from our analytical results, all mutant strategies yield the same average payoff (whereas the payoff of the resident may vary). In two of the three cases, the respective payoff is optimal (A, C). Only in the stag hunt game, equalizers cannot enforce the socially optimal payoff (B). But also in those games (with $T < R$), the socially optimal payoff of mutual cooperation can still be achieved in equilibrium. Players merely need to use the strategy of Always Cooperating instead. For reference, the gray area shows the set of feasible payoffs for the respective game.

computed explicitly. Instead, for our results we focus more on the notion of equalizer strategies, as introduced by Boerljist et al. (47) and Press and Dyson (48). These strategies have the remarkable property that they can unilaterally control the coplayer's payoff, independent of the coplayer's strategy (Fig. 2G). This makes them extremely useful tools to construct Nash equilibria. Once every population member adopts an equalizer strategy, no single player has an incentive to deviate. Interestingly, however, deviators suffer no harm either. In particular, Nash equilibria based on equalizer strategies do not satisfy the stronger notion of evolutionary stability (62).

Evolutionary stability is generally difficult to achieve in repeated games. In fact, for the standard case of infinitely repeated games without errors, no strategy is evolutionarily stable (63–65): one can always identify mutant strategies that may invade by neutral drift. But even identifying strategies that satisfy the weaker condition of being a Nash equilibrium has been difficult in the field of indirect reciprocity. These difficulties are particularly apparent in models with "private information" (21). In such models, individuals may hold different views on which reputation they assign to others. These disagreements may accumulate over time, which makes cooperation difficult to sustain (52, 66, 67). To make analytical progress, the concept of equalizer strategies is particularly convenient. These strategies allow us to give a proof of principle: We rigorously show the existence of cooperative Nash equilibria, for any pairwise social dilemma, for all sufficiently large continuation probabilities, for direct and indirect reciprocity—even under private information.

Interestingly, however, even the most favorable equalizer strategies do not necessarily produce the socially optimal outcome. One counterexample is the stag hunt game (Fig. 5B). Here, equalizers exist, but they cannot ensure the optimal payoff of R. This insufficiency, however, does not diminish our results, nor is it a surprise. Because the stag hunt game's payoffs satisfy both $R > T$ and $R > P$, a coplayer can always avoid an average payoff of R by defecting in all rounds. Hence, unilaterally imposing a guaranteed payoff of R on the coplayer is clearly infeasible. Nevertheless, our more general result, that the game allows for full cooperation in equilibrium, holds. In this game, players simply need to adopt the strategy of always cooperating, instead of adopting an equalizer strategy.

To sum up, social dilemmas are at the core of many collective action problems. To resolve them, people frequently respond to an opponent's previous behavior. Prior to making a decision, they form opinions about their opponent, either based on direct experiences, an opponent's third-party interactions, or both (31, 32). In our work, we mathematically characterize strategies people can use to sustain cooperation, independently of the kind of reciprocity they adopt, and independently of the specific social dilemma at hand.

Methods

Game dynamics and resulting payoffs

We consider a population of n players. Each round, two players are selected uniformly at random. They each play action A or action B and receive a payoff determined by the corresponding entries of the payoff matrix. The $n-2$ players who were not selected for that round receive a payoff of zero. Let $\pi_i(t)$ denote Player i's resulting expected payoff in round t. We define a player's total payoff as the sum of these one-round payoffs times a normalization factor of $(1-d)n/2$, so that the total expected payoff is

$$\pi_i = (1-d)\frac{n}{2}\sum_{t=0}^{\infty} d^t \pi_i(t). \tag{4}$$

The normalization factor ensures that the resulting values are in the same range as the game's one-round payoffs. For example, when all players use action A in all their interactions, the above formula guarantees that each player's expected total payoff (across all interactions and rounds) is R. Alternatively, one may also interpret the above payoff formula to represent a scenario in which individuals discount future payoffs by a constant rate d (68–72). In contrast, Press and Dyson (48), among many other works, consider the asymptotic behavior of a game without discounting. This enables them to compute payoffs by calculating the stationary distribution of the Markov Chain defined by the game process. In the limit of $d \to 1$, our payoff definition approaches theirs.

Except for the case of a donation game among players with reactive strategies (42), there is no known closed-form solution to compute a player's expected payoff $\pi_i(t)$ in a given round t. Hence, for practical purposes, the value of (4) needs to be approximated numerically with simulations. For example, for a given population composition, we may independently simulate the above game dynamics k times. For each run, we sum up each player's payoff across all rounds. Then we sum up these total payoffs across all simulation runs, divide by k, and multiply with the normalization constant $(1-d)n/2$.

In Fig. 5, we have run $k = 5 \times 10^5$ independent simulations for each mutant strategy. For graphical purposes, there we only report payoffs from interactions between the mutant and one given resident, Resident 1 (while ignoring the payoffs from all other interactions). In this case, the relevant normalization constant is $(1-d)n(n-1)/2$. For the mutant, this procedure gives the same result as (4). However, for the resident, the result is different from (4), because we neglect the resident's payoff against other residents. For each of the three panels, the resident strategy is determined by Algorithm 1. All have $\lambda = 0.5$ and $p_0 = 1$. The other entries $(p_{AA}, p_{AB}, p_{BA}, p_{BB})$ are (1.000, 0.331, 1.000, 0.666) for the prisoner's dilemma, (0.498, 1.000, 1.000, 0.498) for stag hunt, and (1.000, 0.498, 0.498, 1.000) for the hawk-dove game.

Instability of nondegenerate reactive strategies

In the Results section, we have argued that in general social dilemmas, only degenerate reactive strategies can be stable. Here, we outline the respective proof (all details are in the Supplementary Information). To this end, consider a resident population of $n > 2$ players, who all adopt the same reactive strategy $\sigma = (p_0, p_A, p_B, \lambda)$. Assume the strategy is nondegenerate, $p_A \neq p_B$, $\lambda < 1$, and that errors are possible, $\varepsilon > 0$. For the proof, we construct a set of four (nonreactive) strategies. Then, we show that at least one of them can invade the resident population.

First, we construct events E_A and E_B. Both E_A and E_B completely define which players are selected and what actions they play in the first $n+3$ rounds, and do so identically apart from the action of player 1 in round 2. The below table summarizes these first $n+3$ rounds. The mutant player is player 1. In round 2, x is action A in E_A and action B in E_B, whereas for y, any consistent choice is permissible. In round 3, $\bar{y}$ is the action that is not y. The dashes indicate actions that are defined by the event, but not specified explicitly in our proof. We show in Supplementary Information Proposition 10 that we can make these choices in such a way that E_A and E_B occur with positive probability.

The invader strategy σ' normally plays in the same way as strategy σ. Only when event E_x has occurred does σ' deviate from σ. It does so by, in one case of $x \in \{A, B\}$, slightly increasing its probability to play action A towards Player 2 next time they are selected to play together, and slightly decreasing it by an identical amount in the other case. Other than that, σ' continues to play exactly like σ.

Construction of equalizer strategies

In the following, we outline how equalizer strategies can be constructed within the space of memory-1 strategies. In the Supplementary Information, we show that for a strategy with cooperation probabilities $p = (p_{AA}, p_{AB}, p_{BA}, p_{BB})^{\intercal}$ and indirectness λ to be a generic equalizer, it needs to have the form

$$\begin{pmatrix} p_{AA} \\ p_{AB} \\ p_{BA} \\ p_{BB} \end{pmatrix} = (\delta^{-1} + \lambda(n-2))(I + \lambda(n-2)M_\varepsilon)^{-1} \begin{pmatrix} 1-\alpha R-\beta \\ 1-\alpha T-\beta \\ -\alpha S-\beta \\ -\alpha P-\beta \end{pmatrix}. \qquad (5)$$

Here, α and β are constants, I denotes the identity matrix and M_ε is the error matrix

$$M_\varepsilon = \begin{pmatrix} 1-\varepsilon & \varepsilon \\ \varepsilon & 1-\varepsilon \end{pmatrix}^{\otimes 2} = \begin{pmatrix} (1-\varepsilon)^2 & (1-\varepsilon)\varepsilon & \varepsilon(1-\varepsilon) & \varepsilon^2 \\ (1-\varepsilon)\varepsilon & (1-\varepsilon)^2 & \varepsilon^2 & \varepsilon(1-\varepsilon) \\ \varepsilon(1-\varepsilon) & \varepsilon^2 & (1-\varepsilon)^2 & (1-\varepsilon)\varepsilon \\ \varepsilon^2 & \varepsilon(1-\varepsilon) & (1-\varepsilon)\varepsilon & (1-\varepsilon)^2 \end{pmatrix}.$$

The payoff that this equalizer strategy enforces is

$$\pi = \alpha^{-1}\left(\frac{1-\delta}{1+(n-2)\delta\lambda}p_0 - \beta\right).$$

Note that whether or not a strategy is an equalizer does not depend on p_0. This general result captures the results of several previous studies as special cases.

1. In Schmid et al. (42), the authors consider the special case of reactive strategies in additive games (with $T \neq P$). Their Supplementary Information Eq. (13) states that a reactive strategy (p_0, p_A, p_B, λ) is an equalizer if and only if

$$p_A - p_B = \frac{1+(n-2)\delta\lambda}{1+(n-2)(1-2\varepsilon)\lambda} \cdot \frac{P-S}{\delta(T-P)}. \qquad (6)$$

   We can recover this result from our (5). For reactive strategies, it takes the form

$$\begin{aligned} \left(\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} + \lambda(n-2)\begin{pmatrix} 1-\varepsilon & \varepsilon \\ \varepsilon & 1-\varepsilon \end{pmatrix}\right)\begin{pmatrix} p_A \\ p_B \end{pmatrix} \\ = (\delta^{-1} + \lambda(n-2)) \cdot \begin{pmatrix} -\alpha S-\beta \\ -\alpha P-\beta \end{pmatrix}, \end{aligned} \qquad (7)$$

   where $\alpha = (T-P)^{-1}$. This is satisfiable (with exactly one β) if and only if condition (6) holds.
2. According to Press and Dyson (48), a memory-1 strategy is an equalizer in an infinitely repeated $(d = 1)$ two-player game $(n = 2)$ if and only if there are constants β (not identical with our β) and γ such that

$$\begin{pmatrix} -1+p_{AA} \\ -1+p_{AB} \\ p_{BA} \\ p_{BB} \end{pmatrix} = \beta\begin{pmatrix} R \\ S \\ T \\ P \end{pmatrix} + \gamma\begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \end{pmatrix} \qquad (8)$$

   This exactly corresponds to our (5) with $\delta = 1$ and $n = 2$ (or alternatively $\lambda = 0$).

Round	**0**		**1**		**2**		**3**		**4**		**5**		**6**		…	**n + 2**	
Players	2	3	1	3	1	2	2	3	1	3	1	3	1	4		1	n
Actions	–	–	–	–	x	y	$\bar{y}$	–	A	–	B	–	–	–		–	–

Algorithm 1 Algorithm to construct an equalizer that enforces the highest possible payoff.

function ComputeEQualizer($R, S, T, P, n, \varepsilon, \lambda$)
  $p_0 \leftarrow 1$
  $\eta \leftarrow \frac{\delta(1+(n-2)\lambda)}{1+(n-2)\delta\lambda}$
  $\zeta \leftarrow \frac{1+\lambda(n-2)}{1+\lambda(n-2)(1-2\varepsilon)}$
  $M_\epsilon \leftarrow$ ErrorMatrix (ϵ)
  $H_\varepsilon \leftarrow (1+\lambda(n-2))^{-1}(I+\lambda(n-2)M_\varepsilon)$
  $(R_\varepsilon, S_\varepsilon, T_\varepsilon, P_\varepsilon)^\intercal \leftarrow H_\varepsilon^{-1}(R, S, T, P)^\intercal$
  if $P_\varepsilon \ / {=} R_\varepsilon$ **then**
    return null
  end if
  $\omega \leftarrow \max\left\{\begin{array}{c} \max\{|R_\varepsilon - T_\varepsilon|, |S_\varepsilon - P_\varepsilon|\}/\eta, \\ (\max\{R_\varepsilon, T_\varepsilon\} - \min\{S_\varepsilon, P_\varepsilon\})/(\eta+\zeta), \\ (\min\{R_\varepsilon, T_\varepsilon\} - \min\{S_\varepsilon, P_\varepsilon\})/\zeta \end{array}\right\}$
  if $\omega > (\min\{R_\varepsilon, T_\varepsilon\} - \max\{S_\varepsilon, P_\varepsilon\})/(\zeta - \eta)$ **then**
    return null
  end if
  $\alpha \leftarrow \omega^{-1}$
  $p_{AA} \leftarrow 1 - \eta^{-1}\alpha(R_\varepsilon - \min\{R_\varepsilon, T_\varepsilon\})$
  $p_{AB} \leftarrow 1 - \eta^{-1}\alpha(T_\varepsilon - \min\{R_\varepsilon, T_\varepsilon\})$
  $p_{BA} \leftarrow 1 - \eta^{-1}\zeta - \eta^{-1}\alpha(S_\varepsilon - \min\{R_\varepsilon, T_\varepsilon\})$
  $p_{BB} \leftarrow 1 - \eta^{-1}\zeta - \eta^{-1}\alpha(P_\varepsilon - \min\{R_\varepsilon, T_\varepsilon\})$
  $\pi \leftarrow \min\{R_\varepsilon, T_\varepsilon\} - (\zeta - 1)/(2\alpha)$
  return $p_0, p_{AA}, p_{AB}, p_{BA}, p_{BB}, \lambda, \pi$
end function

Based on (5), we can also provide an algorithm that takes the game parameters and an indirectness value λ as an input to compute an equalizer strategy that enforces the highest payoff (among all equalizers with indirectness λ). It returns the strategy parameters as well as the payoff π that this strategy enforces, or **null** if no such equalizer strategy exists. The correctness of this Algorithm 1 is shown in Supplementary Information Theorem 6.

Figure parameters

Figure 5: Each of $N = 100$ points represents the payoffs of a randomly generated mutant strategy against $n-1$ identical equalizer residents in a game with $n = 50$ players. Prisoner's dilemma: $R = 3$, $S = 0$, $T = 5$, $P = 1$. Stag hunt: $R = 3$, $S = 0$, $T = 2$, $P = 1$. Hawk-dove: $R = 2$, $S = 0$, $T = 4$, $P = -2$. $\varepsilon = 10^{-3}$, $\delta = 0.99$. Each point is an average over 5×10^5 samples with the same mutant. Equalizer strategies were generated with an implementation of Algorithm 1.

Supplementary Material

Supplementary material is available at *PNAS Nexus* online.

Funding

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.).

Author Contributions

All authors conceived and discussed the study; V.H. analyzed the model; C.H., V.H., and K.C. wrote the main text manuscript; V.H. wrote the Supplementary Information text and conducted simulations; all authors discussed the results and edited both texts.

Data Availability

The data shown in Fig. 5 were generated in a computer simulation written in Rust (compiled with rustc 1.78) and Python 3.10. The computer code and the obtained simulation data are available at https://zenodo.org/doi/10.5281/zenodo.14510769.

References

1 Barfuss W, Donges JF, Vasconcelos VV, Kurths J, Levin SA. 2020. 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(23):12915–12922.
2 Hardin G. 1968. The tragedy of the commons. *Science*. 162(3859): 1243–1248.
3 Kleshnina M, Hilbe C, Šimsa S, Chatterjee K, Nowak MA. 2023. The effect of environmental information on evolution of cooperation in stochastic games. *Nat Commun*. 14(1):4153.
4 Olson M. *The logic of collective action: public goods and the theory of groups*. Harvard University Press, 1971.
5 Broom M, Rychtář J. *Game-theoretical models in biology*. Chapman and Hall/CRC, 2013.
6 Fudenberg D, Tirole J. *Game theory*. 6th ed. Cambridge: MIT Press, 1998.
7 Sigmund K. *The calculus of selfishness*. Princeton University Press, 2010.
8 Nowak MA. 2006. Five rules for the evolution of cooperation. *Science*. 314(5805):1560–1563.
9 Axelrod R, Hamilton WD. 1981. The evolution of cooperation. *Science*. 211(4489):1390–1396.
10 Glynatsi NE, Knight VA. 2021. A bibliometric study of research topics, collaboration and centrality in the field of the Iterated Prisoner's Dilemma. *Humanit Soc Sci Commun*. 8(1):45.
11 Hilbe C, Chatterjee K, Nowak MA. 2018. Partners and rivals in direct reciprocity. *Nat Hum Behav*. 2(7):469–477.
12 Rossetti CSL, Hilbe C. 2023. Direct reciprocity among humans. *Ethology*. 130(4):e13407.
13 Akin E. The iterated Prisoner's Dilemma: good strategies and their dynamics. In: De Gruyter, June 2016. p. 77–107.
14 Hauert C, Schuster HG. 1997. Effects of increasing the number of players and memory size in the iterated prisoner's dilemma: a numerical approach. *Proc Biol Sci*. 264(1381):513–519.
15 Stewart AJ, Plotkin JB. 2013. From extortion to generosity, evolution in the iterated prisoner's dilemma. *Proc Natl Acad Sci U S A*. 110(38):15348–15353.
16 Szolnoki A, Perc M. 2014. Evolution of extortion in structured populations. *Phys Rev E*. 89(2):022804.
17 Martinez-Vaquero LA, Cuesta JA, Sánchez A. 2012. Generosity pays in the presence of direct reciprocity: a comprehensive study of 2×2 repeated games. *PLoS One*. 7(4):e35135.
18 Stewart AJ, Plotkin JB. 2014. Collapse of cooperation in evolving games. *Proc Natl Acad Sci U S A*. 111(49):17558–17563.
19 Nowak MA, Sigmund K. 1998. Evolution of indirect reciprocity by image scoring. *Nature*. 393(6685):573–577.
20 Nowak MA, Sigmund K. 2005. Evolution of indirect reciprocity. *Nature*. 437(7063):1291–1298.
21 Okada I. 2020. A review of theoretical studies on indirect reciprocity. *Games*. 11(3):27.
22 Sigmund K. 2012. Moral assessment in indirect reciprocity. *J Theor Biol*. 299(11):25–30.
23 Clark D, Fudenberg D, Wolitzky A. 2020. Indirect reciprocity with simple records. *Proc Natl Acad Sci U S A*. 117(21): 11344–11349.

24 Nakamura M, Ohtsuki H. 2014. Indirect reciprocity in three types of social dilemmas. *J Theor Biol*. 355:117–127.
25 Murase Y, Hilbe C. 2023. Indirect reciprocity with stochastic and dual reputation updates. *PLoS Comput Biol*. 19(7):e1011271.
26 Ohtsuki H, Iwasa Y. 2004. How should we define goodness?—Reputation dynamics in indirect reciprocity. *J Theor Biol*. 231(1): 107–120.
27 Fischbacher U, Gächter S, Fehr E. 2001. Are people conditionally cooperative? Evidence from a public goods experiment. *Econ Lett*. 71(3):397–404.
28 Montero-Porras E, Grujić J, Fernández Domingos E, Lenaerts T. 2022. Inferring strategies from observations in long iterated prisoner's dilemma experiments. *Sci Rep*. 12(1):7589.
29 Okada I, Yamamoto H, Sato Y, Uchida S, Sasaki T. 2018. Experimental evidence of selective inattention in reputation-based cooperation. *Sci Rep*. 8(1):14813.
30 Wedekind C, Milinski M. 2000. Cooperation through image scoring in humans. *Science*. 288(5467):850–852.
31 Melamed D, Simpson B, Abernathy J. 2020. The robustness of reciprocity: experimental evidence that each form of reciprocity is robust to the presence of other forms of reciprocity. *Sci Adv*. 6(23):eaba0504.
32 Molleman L, Broek E, Egas M. 2013. Personal experience and reputation interact in human decisions to help reciprocally. *Proc Biol Sci*. 280(1757):20123044.
33 Berger U, Grüne A. 2016. On the stability of cooperation under indirect reciprocity with first-order information. *Games Econ Behav*. 98(1435):19–33.
34 Brandt H, Sigmund K. 2004. The logic of reprobation: assessment and action rules for indirect reciprocation. *J Theor Biol*. 231(4): 475–486.
35 Leimar O, Hammerstein P. 2001. Evolution of cooperation through indirect reciprocity. *Proc Biol Sci*. 268(1468):745–753.
36 Michel-Mata S, *et al.* 2024. The evolution of private reputations in information-abundant landscapes. *Nature*. 634(8035):883–889.
37 Radzvilavicius AL, Stewart AJ, Plotkin JB. 2019. Evolution of empathetic moral evaluation. *Elife*. 8:e44269.
38 Santos FP, Santos FC, Pacheco JM. 2018. Social norm complexity and past reputations in the evolution of cooperation. *Nature*. 555(7695):242–245.
39 Pal S, Hilbe C, Glynatsi NE. 2024. The co-evolution of direct, indirect, and generalized reciprocity, arXiv, arXiv:2411.03488. https://doi.org/10.48550/arXiv.2411.03488.
40 Roberts G. 2007. Evolution of direct and indirect reciprocity. *Proc Biol Sci*. 275:173–179.
41 Sasaki T, Uchida S, Okada I, Yamamoto H. 2024. The evolution of cooperation and diversity under integrated indirect reciprocity. *Games*. 15(2):15.
42 Schmid L, Chatterjee K, Hilbe C, Nowak MA. 2021. A unified framework of direct and indirect reciprocity. *Nat Hum Behav*. 5(10):1292–1302.
43 Pacheco JM, Santos FC, Souza MO, Skyrms B. 2009. Evolutionary dynamics of collective action in n-person stag hunt dilemmas. *Proc Biol Sci*. 276(1655):315–321.
44 Skyrms B. *The stag-hunt game and the evolution of social structure*. Cambridge: Cambridge University Press, 2003.
45 Diekmann A. 1985. Volunteer's dilemma. *J Confl Resolut*. 29(4): 605–610.
46 Doebeli M, Hauert C. 2005. Models of cooperation based on the prisoner's dilemma and the snowdrift game. *Ecol Lett*. 8(7): 748–766.
47 Boerlijst MC, Nowak MA, Sigmund K. 1997. Equal pay for all prisoners. *Am Math Mon*. 104(4):303–305.
48 Press WH, Dyson FD. 2012. Iterated prisoner's dilemma contains strategies that dominate any evolutionary opponent. *Proc Natl Acad Sci U S A*. 109(26):10409–10413.
49 Molander P. 1985. The optimal level of generosity in a selfish, uncertain environment. *J Confl Resolut*. 29(4):611–618.
50 Nowak MA, Sigmund K. 1992. Tit for tat in heterogeneous populations. *Nature*. 355(6357):250–253.
51 Glynatsi NE, Akin E, Nowak MA, Hilbe C. 2024. Conditional cooperation with longer memory. *Proc Natl Acad Sci U S A*. 121(50): e2420125121.
52 Hilbe C, Schmid L, Tkadlec J, Chatterjee K, Nowak MA. 2018. Indirect reciprocity with private, noisy, and incomplete information. *Proc Natl Acad Sci U S A*. 115(48):12241–12246.
53 Kerr B, Godfrey-Smith P, Feldman MW. 2004. What is altruism? *Trends Ecol Evol*. 19(3):135–140.
54 Couto M, Pal S. 2023. Introspection dynamics in asymmetric multiplayer games. *Dyn Games Appl*. 13(4):1256–1285.
55 McAvoy A, Rao A, Hauert C. 2021. Intriguing effects of selection intensity on the evolution of prosocial behaviors. *PLoS Comput Biol*. 17(11):e1009611.
56 Nowak MA, Sigmund K. 1998. The dynamics of indirect reciprocity. *J Theor Biol*. 194(4):561–574.
57 Nowak MA, Sigmund K. 1993. A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner's Dilemma game. *Nature*. 364(6432):56–58.
58 Hilbe C, Traulsen A, Sigmund K. 2015. Partners or rivals? Strategies for the iterated prisoner's dilemma. *Games Econ Behav*. 92:41–52.
59 Hilbe C, Nowak MA, Traulsen A. 2013. Adaptive dynamics of exortion and compliance. *PLoS One*. 8(11):e77886.
60 Rand DG, Nowak MA. 2012. Human cooperation. *Trends Cogn Sci*. 117:413–425.
61 Schmid L, Hilbe C, Chatterjee K, Nowak MA. 2022. Direct reciprocity between individuals that use different strategy spaces. *PLoS Comput Biol*. 18(6):e1010149.
62 Maynard Smith J, Price GR. 1973. The logic of animal conflict. *Nature*. 246(5427):15–18.
63 Boyd R, Lorberbaum JP. 1987. No pure strategy is evolutionarily stable in the repeated prisoner's dilemma game. *Nature*. 327(6117):58–59.
64 García J, Veelen M. 2018. No strategy can win in the repeated prisoner's dilemma: linking game theory and computer simulations. *Front Robot AI*. 5:102.
65 Lorberbaum JMD. 1994. No strategy is evolutionary stable in the repeated Prisoner's Dilemma. *J Theor Biol*. 168(2):117–130.
66 Murase Y, Hilbe C. 2024. Computational evolution of social norms in well-mixed and group-structured populations. *Proc Natl Acad Sci U S A*. 121(33):e2406885121.
67 Uchida S. 2010. Effect of private information on indirect reciprocity. *Phys Rev E Stat Nonlin Soft Matter Phys*. 82(3):036111.
68 Abreu D. 1988. On the theory of infinitely repeated games with discounting. *Econometrica*. 56(2):383–396.
69 Fudenberg D, Maskin E. 1986. The folk theorem in repeated games with discounting or with incomplete information. *Econometrica*. 54(3):533–554.
70 Ichinose G, Masuda N. 2018. Zero-determinant strategies in finitely repeated games. *J Theor Biol*. 438:61–77.
71 McAvoy A, Hauert C. 2016. Autocratic strategies for iterated games with arbitrary action spaces. *Proc Natl Acad Sci U S A*. 113(13):3573–3578.
72 Nowak MA, Sigmund K. 1995. Invasion dynamics of the finitely repeated prisoner's dilemma. *Games Econ Behav*. 11(2):364–390.