



Payoff equivalence and complete strategy spaces of direct reciprocity

Philip LaPorte^{a,b,1} , Christian Hilbe^{c,d} , Nikoleta E. Glynatsi^{c,e} , and Martin A. Nowak^{b,f}

Affiliations are included on p. 8.

Edited by Marcus Feldman, Stanford University, Stanford, CA; received July 10, 2025; accepted December 24, 2025

Repeated games and stochastic games are important frameworks to study direct reciprocity. Individuals react strategically to their coplayers' previous behavior. While strategies in such games can be arbitrarily complex, explorations of evolutionary dynamics are often done in specific strategy spaces. Individuals may consider a fixed number of past rounds, or only some of the partner's previous actions. Such restrictions can make the interpretation of the results difficult. Strategies found to be superior within a restricted set may lose stability when more complex strategies are permitted. We discuss two notions of completeness that rule out this possibility. If a strategy space, S , is best-reply-complete, then any strategy in S is guaranteed to have a best reply in S . If a space, S , is payoff-complete, then any strategy playing against an opponent in S can be replaced by an equivalent strategy within S without affecting either player's payoff. Sufficient conditions for best-reply-completeness have been given in a seminal paper by Levínský et al. Here, we show that for strategies of bounded memory, the same conditions are sufficient for payoff-completeness. Furthermore, using those conditions, we illustrate how to construct many complete spaces for simple games. Taken together, our findings highlight the importance of complete strategy spaces, which are particularly useful when interpreting evolutionary simulations and determining best responses.

evolutionary game theory | repeated prisoner's dilemma | direct reciprocity | stochastic processes | cooperation

Direct reciprocity is an important mechanism for cooperation (1). This mechanism focuses on individuals who interact for many rounds, and hence are able to react to their coplayer's previous actions. If players engage in the same stage game in every round, researchers speak of a repeated game (2). The repeated prisoner's dilemma (3–10) is a classic and well studied example. Games like this help us capture the logic of reciprocal cooperation in animals (11), humans (12), and across nations (13). More generally, when the stage games played in each round can change, depending on the previous game and the players' previous actions, researchers speak of a stochastic game (14). Here, the stage game is associated with a dynamic *game state*, which can transition from round to round (Fig. 1A). The game state affects the actions available to players and the payoffs which they earn.

To make their decisions in each round, individuals follow their strategies. A strategy is a rule that tells the player what to do, given the history of previous play and the current (game) state. In general, such strategies can be rather complex. For computational reasons, however, evolutionary analyses of games often assume that players choose from a restricted set of strategies. Well-known sets are reactive strategies (15, 16), memory-1 strategies (17), or memory- n strategies (18–22). These sets describe individuals who hold in memory only the coplayer's last action, the outcome of the last round, or the outcome of the last n rounds, respectively. While restricting memory to the most recent experiences may seem innocuous, results from respective studies need to be interpreted with caution. A specific strategy might be successful within a restricted strategy space, but once we allow for strategies of different complexity it might be outperformed. As a result, it can be difficult to ensure that a winning strategy derived from evolutionary simulations is robust in a more general environment. To ascertain robustness, it becomes important to investigate the following broad question: Can strategies outside of a given space achieve superior outcomes which are impossible within the space?

A key contribution toward that question was made by Levínský et al. (23), a work we briefly sketch below and that we revisit in more detail in *Results*. Informally, they give a condition under which a strategy of a certain kind has a best reply of the same kind. To state their result, Levínský et al. introduce "factors." A factored strategy (23) chooses

Significance

When people interact repeatedly, they often adopt strategies of direct reciprocity. In principle, such strategies can be very complex. A person's next move could depend on every prior interaction. Yet, many models restrict the set of feasible strategies, for example, by limiting the number of past events individuals hold in memory. Such restrictions come with a risk: Conclusions for restricted strategy spaces may no longer apply if individuals were allowed to deviate toward more complex strategies. A strategy space is called complete if that risk is provably nonexistent. Here, we introduce two complementary notions of completeness. Taking an established criterion for one of these notions, we extend it to the other notion and construct many complete spaces.

Author contributions: P.L., C.H., N.E.G., and M.A.N. designed research; performed research; and wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

Copyright © 2026 the Author(s). Published by PNAS. This article is distributed under Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 (CC BY-NC-ND).

¹To whom correspondence may be addressed. Email: philip_laporte@berkeley.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2518486123/-DCSupplemental>.

Published January 29, 2026.

actions based not on the game history directly, but rather on some simple function φ of the history. This function φ is the factor. A factor is like a lens through which a strategy views the game history (Fig. 1 *B* and *C*). It filters out some information and registers only a desired portion, such as a given player's previous two moves. That portion alone informs the focal player's next action. One important property a factor may have is *recursiveness*. Recursiveness means that the information stored by a factor can be updated incrementally (Fig. 2). This prevents a factor from being able to store more information about less recent rounds than about more recent rounds. For stochastic games with a fixed discount factor, Levínský et al. (23) prove the following result: Let φ be a recursive factor. If σ_1 is a φ -factored strategy for player 1, then there is also a (deterministic) φ -factored strategy σ_2 for player 2, which is a best reply to σ_1 . Furthermore, the best reply property can be guaranteed in all subgames. This result means that to any strategy, there is a best reply that takes into account the same kind of information. The above result can be used to identify strategy spaces that are closed under taking best replies. We propose to call such strategy spaces *best-reply-complete*.

Understanding the scope of best replies is important for carrying out informed evolutionary studies. However, for many questions in evolutionary game theory, best replies are not the only relevant concept. For example, if the strategy space is uncountably large, random mutations typically do not introduce best replies into the population; however, they may well introduce *better replies* (strategies that yield a better payoff than the payoff of a resident strategy). Furthermore, in evolutionary competitions between two strategies, typically the payoff of both types matter. To take into account such considerations, we propose to consider strategy spaces that are *payoff-complete*. To this end, fix a focal player's strategy space and suppose the focal player adopts an element of that space, whereas the coplayer adopts an arbitrary strategy. Then payoff-completeness means that the coplayer can always switch to a strategy in the focal player's space without affecting either player's payoff. A property of this kind was described for the space of memory- n strategies by Press and Dyson in the case of no discounting (24). This property ensures, for example, that a focal player with a zero-determinant strategy enforces a linear payoff relationship against all possible coplayers, not only against those players with strategies of similar complexity (24–30).

To give a criterion for a strategy space to be payoff complete, we ask a question analogous to the one by Levínský et al. (23): If an arbitrary strategy is used against a fixed φ -factored opponent, can this arbitrary strategy be exactly emulated—in terms of the payoffs to both players—by an alternative φ -factored strategy? We address this question for stochastic games with a fixed discount factor and finitely many states. Under appropriate assumptions, we show if σ_1 is a φ -factored strategy for player 1, and σ_2 is an arbitrary strategy for player 2, then there is an alternative φ -factored strategy σ'_2 for player 2 that results in the same payoffs to both players (Fig. 3). Furthermore, this payoff equivalence can be guaranteed for all payoff functions simultaneously (i.e., σ'_2 does not depend on the specific entries of the payoff matrix).

Our results, when combined with those of Levínský et al. (23), imply that many diverse strategy spaces are both best-reply-complete and payoff-complete. Complete spaces are essential for carrying out evolutionary studies with robust conclusions. We illustrate how to easily identify examples of complete spaces using partitions. We also point out the incompleteness of other strategy spaces.

Results

Our following analysis consists of two parts. First, we formally introduce stochastic games, factors, and recursiveness. Here, we also explain the results of Levínský et al. (23) in order to contextualize our own main result, which is presented immediately afterward. In the second part, we illustrate these concepts for the example of a repeated prisoner's dilemma, and we define the notion of complete strategy spaces. Throughout this main text, we provide a summary of our results; all proofs and more detailed information are in *SI Appendix*.

Stochastic Games. Stochastic games are a model for repeated interactions in dynamic environments. A stochastic game (Fig. 1*A*) is played over infinitely many rounds. In each round t , there is a game state z_t which belongs to a set of possible states S . Player 1 chooses an action a_t^1 from a set $A_1(z_t)$ of available actions in that state. Likewise player 2 chooses an action a_t^2 from $A_2(z_t)$. The result is an action profile $a_t = (a_t^1, a_t^2)$. At the end of each round, each player i is assigned an immediate payoff $u_i(z_t, a_t)$ based on the game state and actions taken. The game state then transitions from state z_t to state z_{t+1} according to a transition probability which depends on the actions a_t .

In general, the players' actions may depend on what happened in all previous rounds (i.e., on the game history) (Fig. 1*B*). We write the set of all possible finite game histories as

$$H^\infty = \left\{ (z_1, a_1, z_2, \dots, z_t) \mid t \in \mathbb{N}, z_j \in S, a_j^i \in A_i(z_j) \right\} \quad [1]$$

A strategy σ_i for player i specifies the probabilities for choosing each action, based on the history. In other words, σ_i is a map which takes any history $h \in H^\infty$ as an input, and outputs a probability distribution $\sigma_i(h)$ over the set of available actions $A_i(z_t)$ in the current state. If player 1 uses a strategy σ_1 and player 2 uses a strategy σ_2 , then the β -discounted payoff to player i is given by

$$\pi_\beta^i(\sigma_1, \sigma_2) = \mathbb{E} \left[(1 - \beta) \sum_{t=1}^{\infty} \beta^{t-1} u_i(z_t, a_t) \right] \quad [2]$$

That is, players aim to maximize their expected discounted payoff across all future rounds.

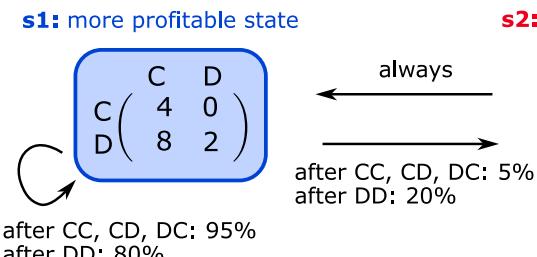
Factors. A factor (23) is a formal device which constrains the complexity of a strategy. The point is that a factored strategy must choose actions based on some simple function of the history, rather than directly based on the game history itself (Fig. 1*C*). Formally, a factor φ is a map from H^∞ to some set X . One thinks of $\varphi(h)$ as some information about h that is stored by φ . We require φ to be such that

$$\varphi(z_1, a_1, \dots, z_t) = \varphi(z'_1, a'_1, \dots, z'_t) \implies z_t = z'_t \quad [3]$$

That is, φ must store the current game state (it distinguishes between two histories h and h' in which the current game state is different). A strategy σ_i for player i is called φ -factored if $\sigma_i(h) = \sigma_i(h')$ whenever $\varphi(h) = \varphi(h')$. In other words, when the factor φ cannot distinguish between two histories, a φ -factored strategy does not distinguish between the two histories either.

Some of the simplest factors are what we call memory- k factors. Such factors only depend on the most recent k rounds. More precisely, if φ is a memory- k factor, and h is a finite history of length at least k , then $\varphi(h)$ is computed by first extracting just the outcome of the most recent k rounds— h_k —and applying some

A A stochastic game with two states



B A game history after five rounds

Round no.	1	2	3	4	5	6
Game state:	s1	s1	s1	s1	s2	s1
Player 1:	C	C	D	D	C	C
Player 2:	C	D	D	D	D	D

C A factored strategy

Round no.	1	2	3	4	5	6
Game state:	s1	s1	s1	s1	s2	s1
Player 1:	C	C	D	D	C	C
Player 2:	C	D	D	D	D	?

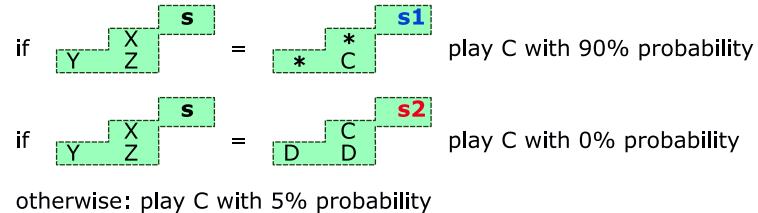


Fig. 1. Stochastic games. A stochastic game involves infinitely many rounds. In each round, there is a game state, which determines the available actions to players and their resulting payoffs. The actions also lead the game state to transition stochastically from one possible state to another. The game state could reflect the external environment, or the actions of additional players who are treated as part of Nature. A stochastic game with one game state is a repeated game. (A), we present an example stochastic game with two states. In both states, the players choose between the actions C (“cooperate”) and D (“defect”) of a basic cooperative dilemma. The immediate payoffs in state 2 are half as large. (B), we show a possible game history after five rounds. (C), we indicate one possible factor and a corresponding factored strategy. The factor isolates the current game state, the actions of the previous round, and player 2’s action in the second-previous round. A factored strategy always decides what to play next based only on this information. In general, factors can be more complicated. For example, they may have unbounded memory, or may be able to distinguish between some actions and not others in the game history.

fixed function φ_k to it. That is, $\varphi(h) = \varphi_k(h_k)$. If h has length less than k , we simply define $\varphi(h) = h$. If φ is a memory- k factor, then a φ -factored strategy is a memory- k strategy (18–22). Such strategies decide what to play based only on the outcome of the most recent k rounds (during the very first k rounds, the strategy’s behavior is not constrained at all. However, see *SI Appendix* for other possibilities).

Recursiveness. Recursiveness (23) is an important property of some factors. It expresses that information about the game history can be updated incrementally. Formally, a factor φ is *recursive* if $\varphi(z_1, a_1, \dots, z_t, a_t, z_{t+1})$ is determined by $\varphi(z_1, a_1, \dots, z_t)$, by a_t , and by z_{t+1} . That is, given the information taken into account in the previous round, the outcome of that round, and the new state, the player has all the information required to make a decision in the next round (Fig. 2). This property prevents a factor from having a greater memory capacity for rounds more distant in the past. Many basic factors are recursive. An example of a non-recursive factor is one for which $\varphi(z_1, a_1, \dots, z_t) = (a_{t-2}, z_t)$. This factor stores only the current game state z_t and the second-previous round of actions, a_{t-2} . The memory cannot be updated incrementally, because the new increment of information a_t, z_{t+1} cannot recover the forgotten round of actions a_{t-1} .

Two Key Theorems on φ -Factored Strategies. Based on these concepts, Levínský et al. prove the following result (theorem 4.1 in ref. 23). Consider a stochastic game with countably many states, finitely many available actions in each state, a bounded payoff function u_2 , and a discount factor $\beta \in (0, 1)$. Let φ be a recursive factor and σ_1 be a φ -factored strategy for player 1. Then there is a deterministic φ -factored strategy σ_2 for player 2, which is a best reply to σ_1 in all subgames. This means that playing σ_2 is optimal, regardless of what moves have already been played. The upshot of this theorem is that if my opponent is recursively φ -factored, I can optimize my strategy by also choosing a φ -factored

strategy. The authors extend this basic result in many directions. In particular, they also consider best replies for all sufficiently large discount factors, and for the limit of no discounting.

In the context of payoff equivalence, we can state our main result as follows. Consider a stochastic game with finitely many states and a discount factor $\beta \in (0, 1)$. Let φ be a recursive factor which is memory- k for some k , let σ_1 be a φ -factored strategy for player 1, and let σ_2 be an arbitrary strategy for player 2, not necessarily even of bounded memory. Then there is a φ -factored strategy σ'_2 for player 2 such that the β -discounted payoffs $\pi_\beta^i(\sigma_1, \sigma'_2)$ for both players are the same as the original payoffs $\pi_\beta^i(\sigma_1, \sigma_2)$. Moreover, these equalities hold regardless of the payoff functions u_i . The upshot of our theorem is that, if my opponent is recursively φ -factored, then any strategy of arbitrary complexity I may wish to adopt has a payoff-equivalent strategy which is also φ -factored.

After having formulated the key concepts and results for φ -factored strategies, in the following, we illustrate their significance in the simple setting of repeated games. We also present a natural motivating application, the identification of complete strategy spaces for direct reciprocity.

Repeated Games. Repeated games can be interpreted as stochastic games with only one static game state. Players meet in infinitely many rounds and choose actions from a fixed finite set A each time, receiving payoffs according to a payoff matrix. For example, in a 2×2 symmetric repeated game such as the repeated prisoner’s dilemma (31–34), each player chooses between two actions, C and D (for *cooperate* and *defect*). They receive payoffs in each round according to a matrix

$$C \quad D \\ C \left(\begin{array}{cc} a & b \\ c & d \end{array} \right). \quad [4]$$

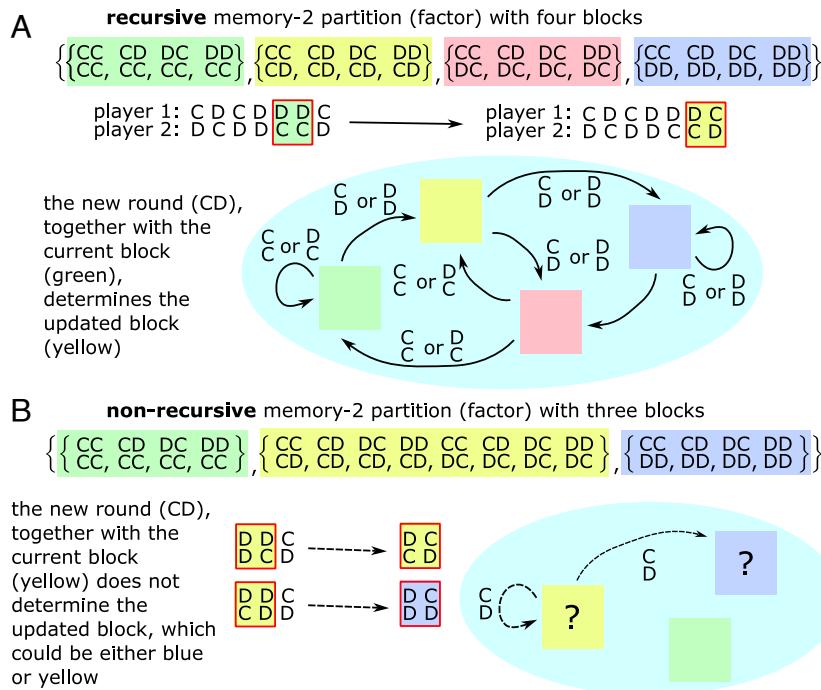


Fig. 2. Recursive and nonrecursive factors. A factor serves to group game histories into blocks, according to their factor values. (A factored strategy then chooses a move based not on the history but on the block.) It can be represented by a partition of histories. Here, we show two memory-2 factors, or partitions, for a simple repeated game with actions C and D. A factor is recursive if the updated block is solely determined by the current block and the new round. Then the game is described by a deterministic finite-state automaton with transitions between the blocks. (A), this memory-2 partition is recursive. It has four blocks, colored green, yellow, red, and blue. Here, the block describing the most recent two rounds is updated from green to yellow when the new actions CD are played. To know that the updated block is yellow, we only need to know that the current block is green and the new actions are CD. (B), this memory-2 partition is nonrecursive. It has three blocks. The blocks represent whether the coplayer has cooperated twice (green), once (yellow), or never (blue) during the last two rounds. Assume the current block is yellow and the outcome of the most recent round is CD. Depending on the precise memory-2 history, the new block could be either yellow or blue. Since the choice is not determined, the partition is not recursive. However, every nonrecursive partition can be made recursive by subdividing the blocks further. Notice that the first partition is a refinement of the second. In the worst-case scenario, one can subdivide into the partition of singletons, which is recursive.

Here, each entry refers to the payoff of the row-player. The respective payoff of the column player follows from symmetry.

Memory- k Factors for Repeated Games. Let H be the set $\{CC, CD, DC, DD\}$ of possible outcomes for a single round. For example, CD means that player 1 played C and player 2 played D. The Cartesian product $H^n = H \times \dots \times H$ is the set of possible outcomes for n consecutive rounds. We write a multiround outcome using “|” notation: For example, CC|CD|DD refers to a sequence of three consecutive rounds. In the most recent round, both players played D. A factor φ then takes all finite game histories and maps them to some arbitrary set X . What a factor really does is allow us to group the histories into “blocks;” each block of histories gets mapped to a unique value $x \in X$. The value x need not have any inherent meaning. In practice, x is just an “atom” which represents the block. We may produce factors simply by constructing appropriate blocks of histories.

For example, there is a memory-1 factor which groups together all histories ending in CC into one block; all histories ending in CD or DC into a second block; and all histories ending in DD into a third block. We represent this factor by a partition of H^1 as follows:

$$P = \{\{CC\}, \{CD, DC\}, \{DD\}\} \quad [5]$$

As another example, there is a memory-1 factor which groups together all histories ending in CC or DC into one block; and all histories ending in CD or DD into a second block. We represent this factor by a partition of H^1 as follows:

$$P = \{\{CC, DC\}, \{CD, DD\}\} \quad [6]$$

As a third example, there is a memory-2 factor which groups together all histories ending in CC|CC into one block; and all histories ending in other two-round patterns, like CC|CD or DD|DD, together into a second block. We represent this factor by a partition of H^2 as follows:

$$P = \{\{CC|CC\}, \{CC|CD, \dots, DD|DD\}\} \quad [7]$$

In this way, every memory- k factor can simply be represented as a partition of H^k . This representation is convenient because each block in P contains just the relevant parts of the histories, that is, the outcomes of the most recent k rounds. Thus an infinite block of many possible game histories is represented by a small, finite block of the partition P (technically, we have ignored histories of length smaller than k . For simplicity, we place each of these in its own unique block of histories, but we do not include them in the partition representation for conciseness).

There are fifteen memory-1 partitions in total (SI Appendix, Fig. S1). Because here we focus on a symmetric game, each partition has a well-defined conjugate partition, which is produced by switching the actions of player 1 and player 2 in each history. For example, the conjugate partition of [6] is $\{\{CC, CD\}, \{DC, DD\}\}$. Some partitions, like the memory-1 partition [5] or the memory-2 partition [7], are equal to their own conjugate. We call these symmetric partitions. The significance of symmetric partitions will be explained below.

Sometimes it is useful to subdivide the blocks of a partition P to get another partition Q with more blocks. In that case, we write $P \leq Q$ (after all, P has fewer blocks) and say Q is a refinement of P . The most refined partition is the partition of H^k into singletons. In that case, the player distinguishes every single memory- k history.

Recursiveness. Suppose P is a memory- k partition. We write $[h]_P$ to represent the partition element or block which contains $h \in H^k$. In terms of partitions, we can interpret the property of recursiveness as follows: For any sequence of outcomes for k rounds $(h^1, \dots, h^k) \in H^k$ and any new outcome $h^{k+1} \in H$, the block for the updated sequence is completely determined by the block for the old sequence, together with the new round. In other words, there is an assignment $P \times H \rightarrow P$ such that

$$(\underbrace{[h^1, \dots, h^k]_P, \quad \text{new round}}_{\text{previous block and}}, \quad \underbrace{h^{k+1}}_{\text{determine}}) \mapsto \overbrace{[h^2, h^3, \dots, h^{k+1}]_P}^{\text{updated block}} \quad [8]$$

Clearly, all memory-1 partitions are recursive. Here, the next moves of the two players fully determine the next block. But only 78,721 memory-2 partitions are recursive (out of about ten billion in total). For an example of such a recursive memory-2 partition (and an example of a partition that is not recursive), see Fig. 2. As n increases, the proportion of recursive partitions among all memory- n partitions shrinks further. But most of the nonrecursive partitions are hard to interpret, and unlikely to be relevant to evolutionary game theory. Moreover, any partition can be refined until it is recursive, as explained in more detail in *SI Appendix*.

Examples of recursive memory- k partitions can be constructed by combining different memory-1 partitions. To see how this works, take a collection of memory-1 partitions P_1, \dots, P_n and define the product partition $P = P_1 \times \dots \times P_n = \{b_1 \times \dots \times b_n \mid b_i \in P_i\}$. The definition simply means that for all $h, h' \in H^k$,

$$[h]_P = [h']_P \iff \forall j: [h^j]_{P_j} = [h'^j]_{P_j} \quad [9]$$

Here, h^j means the outcome of the j th round, in chronological order, in the history h . Such a product partition is recursive if and only if $P_1 \leq P_2 \leq \dots \leq P_n$. That is, the further events are in the past, the coarser the respective partitions tend to be. The intuition is that stored memory about a given round must monotonically decrease as new rounds are played and the round recedes into the past. Some of these product partitions are well known in evolutionary game theory, and we will allude to examples below. However, product partitions are not the only possibilities.

Factored Strategies. A memory- k strategy (18–22) decides the next move based only on the outcome of the previous k rounds. For example, memory-0 strategies choose what to play without regard to the previous rounds. These include the strategy always-cooperate (ALLC) and always-defect (ALLD). Memory-1 strategies choose what to play based only on the previous round. The strategy win-stay, lose-shift (WSLS) (17, 35) cooperates only if both players chose the same action in the previous round. The strategy generous tit-for-tat (GTFT) (15, 36) cooperates if the opponent cooperated in the previous round, and cooperates with some probability q otherwise. Naturally, memory- k factors give rise to subspaces of memory- k strategies. We give two characteristic examples.

For the first example, we consider the (symmetric) partition given by the trivial memory-1 partition $\{\{\text{CC}\}, \{\text{CD}\}, \{\text{DC}\}, \{\text{DD}\}\}$. This partition represents a factor φ that groups histories together into four blocks, according to the exact outcome of the previous round. A φ -factored strategy must specify four probabilities, p_1, p_2, p_3, p_4 . Here, p_1 corresponds to the first block. It is the probability to cooperate if the previous round was CC. Likewise, p_2 corresponds to the second block, and so on. In addition to these conditional cooperation probabilities, the strategy must also give a probability p_0 for cooperating in the very first round. Therefore, in this case, a φ -factored strategy is the same thing as a memory-1 strategy. Now, applying the theorem of Levinský et al. (23) and our theorem, we conclude the following two facts: *i*) Every memory-1 strategy has a best reply (in all subgames) which is also memory-1. *ii*) If the focal player uses a memory-1 strategy and the coplayer uses an arbitrary strategy, the coplayer can switch to a memory-1 strategy while preserving the payoffs to both players [this extends a result of Press and Dyson (24) to the case of discounting.]

For the second example, take the memory-1 partition [6], $\{\{\text{CC, DC}\}, \{\text{CD, DD}\}\}$. This partition represents a factor φ that groups histories together into two blocks, according to player 2's last move. A φ -factored strategy must specify two conditional cooperation probabilities, one for each block. Because the decision of whether or not to cooperate is based on player 2's most recent move, the factor introduces an important asymmetry between the two players. A φ -factored strategy for player 1 responds to the opponent's previous move. It is called a reactive-1 strategy (15, 16). An example is GTFT. On the other hand, a φ -factored strategy for player 2 responds to player 2's own previous move. It is called a self-reactive-1 strategy (37). An example is Alternator, a strategy which always plays the opposite move as it played in the previous round. Now, applying the theorem of Levinský et al. (23) and our own theorem, we conclude the following two facts: *i*) Every reactive-1 strategy has a best reply (in all subgames) which is self-reactive-1. *ii*) If the focal player uses a reactive-1 strategy, and the coplayer uses an arbitrary strategy, then the coplayer can switch to a self-reactive-1 strategy while preserving the payoffs of both players. Again, this example extends a previous result (37) to the case of discounting.

The important difference between the two examples is that the first partition is symmetric. Only in this first example, we recover the same partition if we swap the moves of the two players. In the second example, there are two equivalent ways to look at the relationship between reactive-1 and self-reactive-1 strategies. The first is to consider the factor φ to be fixed and the strategy types to arise for player 1 and player 2, respectively. The second way is to focus on player 1 exclusively and to consider the strategy types to arise from two separate factors, one based on the partition $\{\{\text{CC, DC}\}, \{\text{CD, DD}\}\}$ and the other based on its conjugate. Both pictures are useful for different settings. The first picture connects well to our earlier key theorems. The second picture is advantageous because it does not require tracking the identities of indistinguishable players. Only a focal player, always designated as player 1, is relevant.

Some other examples of the symmetric case include the partitions [5] and [7]. Partition [5] leads to the space of memory-1 counting strategies (19, 38). Counting strategies count how many players cooperated in the previous round and respond accordingly. The k -fold product of this partition leads to the space of memory- k roundwise-counting strategies. Some other notable examples of the asymmetric case include the k -fold product of

[6] or its conjugate. This leads to the reactive- k strategies and self-reactive- k strategies (37, 39).

Overall, we note that memory- k partitions are either symmetric or asymmetric. The asymmetric partitions come in conjugate pairs. For evolutionary simulations, partitions with fewer blocks are often more practical, because they generate strategy spaces with fewer dimensions.

Completeness. In the previous section, we illustrated how the theorem of Levínský et al. (23) and our own theorem have meaningful implications for two different recursive partitions. We saw that the implications depended on whether the partitions were symmetric or asymmetric. In this section, we explain how these theorems establish certain important completeness properties.

For symmetric partitions, the situation is simple. Each symmetric partition gives rise to a single space of strategies. For example, for the partition $\{\{CC\}, \{CD, DC\}, \{DD\}\}$, we obtain the memory-1 counting strategies. That strategy space \mathbf{S} is guaranteed to enjoy two completeness properties. *Best-reply-completeness* means that for each strategy $s \in \mathbf{S}$, there exists a strategy $s' \in \mathbf{S}$ that is a best reply to s in all subgames. This means that given the history of the game up to any point, no strategy can achieve a greater continuation payoff against s than the payoff s' achieves against s . *Payoff-completeness* means that the following holds: for every strategy $s \in \mathbf{S}$ for the focal player, and every strategy s' for the coplayer (of arbitrary complexity), there is a strategy $s'' \in \mathbf{S}$ for the coplayer which achieves the same payoffs to both players. We call this s'' a payoff-preserving reply. Payoff-completeness then guarantees the existence of payoff-preserving replies. Best-reply-completeness guarantees the existence of a best reply.

For asymmetric partitions, the situation is a bit more complex. Each asymmetric partition and its conjugate, such as $\{\{CC, DC\}, \{CD, DD\}\}$ and $\{\{CC, CD\}, \{DC, DD\}\}$, give rise to two distinct strategy spaces—in this case, the spaces of reactive-1 and self-reactive-1 strategies. Neither of these is necessarily complete on its own. However, a best reply or payoff-preserving reply to a strategy in one space can always be found in the other space. In particular, their *union* is best-reply-complete and payoff-complete.

Fig. 3 illustrates payoff-completeness in a repeated prisoner's dilemma. Here, player 1 uses a memory-1 counting strategy $s_1 = (p_{\{CC\}}, p_{\{CD, DC\}}, p_{\{DD\}}) = (1, 1/4, 0)$. Player 2 uses a memory-1 strategy $s_2 = (p_{\{CC\}}, p_{\{CD\}}, p_{\{DC\}}, p_{\{DD\}}) = (3/4, 1, 1/4, 1/2)$. (Note that in our convention, the subscripts CD and DC are ordered with player 1's move first and player 2's move second.) The figure confirms that player 2 can switch to a memory-1 counting strategy $s'_2 = (3/4, 7/16, 1/2)$, without changing the payoff of either player.

Fig. 4 A and B provides a different illustration of the same insight. Again, player 1 uses some fixed memory-1 counting strategy. For player 2 we sample a large number of strategies, either from the space of memory-1 strategies (*Left*) or from the memory-1 counting strategies (*Right*). In line with our results, the two generated payoff spaces coincide.

Fig. 4 C and D illustrates a strategy space which is incomplete in both senses, the space of reactive-2 strategies. In this example, player 1 uses a fixed reactive-2 strategy. Player 2 varies over a sample of other reactive-2 strategies (*Left*) or self-reactive-2 strategies (*Right*). We observe that player 1's strategy is a best reply to itself in the space of reactive-2 strategies, even though it can be invaded by a self-reactive-2 strategy. This is an example where

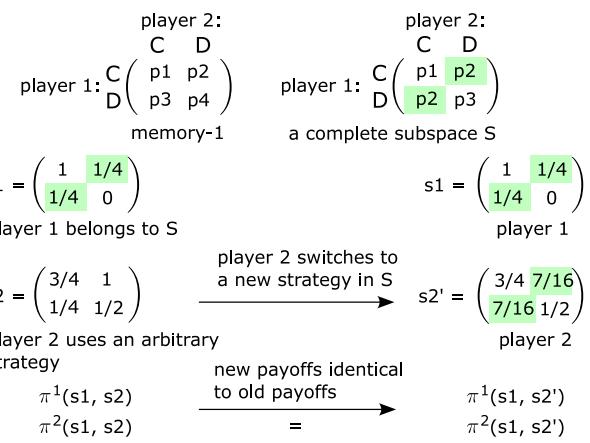


Fig. 3. Payoff-completeness in action. A strategy space \mathbf{S} for a symmetric game is payoff-complete if the following holds: If the focal player uses a strategy in \mathbf{S} and the coplayer uses an arbitrary strategy, then the coplayer can switch to a strategy in \mathbf{S} without changing the payoff of either player. Equivalently: If my opponent has a fixed strategy in \mathbf{S} and I am free to vary my strategy, then I can realize all feasible payoff outcomes by varying my strategy in \mathbf{S} only. Here we show an example of a payoff-complete subspace of memory-1 strategies, the space of memory-1 counting strategies. Memory-1 counting strategies only depend on how many of the two players cooperated the previous round. Suppose player 1 uses a counting strategy $\mathbf{p} = (1, 1/4, 1/4, 0)$ while player 2 uses a memory-1 strategy $\mathbf{q} = (3/4, 1, 1/4, 1/2)$ as shown. Then player 2 can switch to the payoff-equivalent counting strategy $\mathbf{q}' = (3/4, 7/16, 7/16, 1/2)$. Neither player's payoff is affected. This conclusion does not depend on the payoff matrix of the game. While our main results are stated for a discount factor $\beta < 1$, the property of payoff-equivalence also applies in the limit $\beta \rightarrow 1$ when payoffs are uniquely defined. For convenience, we have assumed this limit when computing the new strategy and the payoffs.

evolutionary results for a strategy space can be misleading. When running simulations restricted to reactive-2 strategies, player 1's strategy in Fig. 4 C and D might well arise as a stable outcome. Yet if the simulation also allowed for (equally complex) self-reactive-2 strategies, this outcome would no longer be robust.

We conclude with a few remarks. First, as defined here, payoff-completeness is concerned with finding a strategy $s'' \in \mathbf{S}$ that preserves payoffs for the entire game, given the two players start in the initial round. Instead, one may be interested in payoff-preserving replies starting from arbitrary subgames. While under the appropriate conditions one can indeed find the respective payoff-preserving replies in \mathbf{S} , they might depend on the subgame—on the game history up to the time of the switch. The situation is different for best-reply completeness. Here, under appropriate assumptions, one can find strategies that are a best reply in all subgames (23).

Second, herein, we focus exclusively on the case of discounted games, with $\beta \in (0, 1)$. For the complete spaces that satisfy our criteria, however, Levínský et al. have also established best-reply-completeness in the limiting average payoff case (no discounting). Addressing this limit for payoff-completeness is a possible topic for future work.

Finally, we note that our result provides a sufficient condition for completeness, but not a necessary one. To illustrate this point, consider the space of reactive-1 strategies. Since the associated partition is not symmetric, this space, taken by itself, does not satisfy our criteria for completeness. Yet it turns out that the space is nevertheless payoff-complete for some particular payoff matrices (e.g., a donation game, which is additive 39, Fig. 5A), despite being incomplete for others (e.g., a snowdrift game, Fig. 5B).

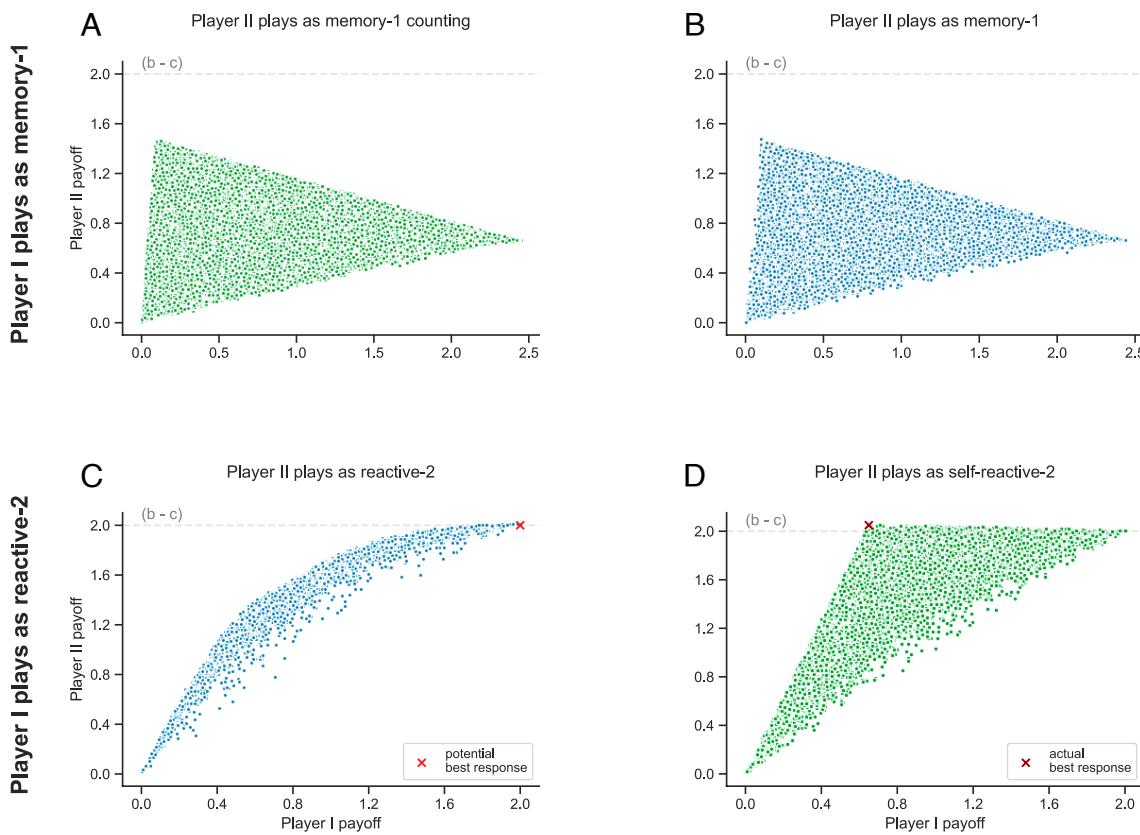


Fig. 4. Complete and incomplete strategy spaces. We examine a prisoner's dilemma with $(a, b, c, d) = (2, -1, 3, 0)$. Although our results are stated for discount rate less than 1, analogous phenomena arise in the limit of no discounting $\beta \rightarrow 1$, which we use here for convenience. (A and B), We fix player 1's strategy to be the memory-1 counting strategy $\mathbf{p} = (0.42, 0.70, 0.70, 0.00)$. For the strategy of player 2, we randomly sample a large number of memory-1 strategies (A) or memory-1 counting strategies (B). The covered area of payoffs is the same in both cases because the space of memory-1 counting strategies is based on a symmetric and recursive partition, and hence is payoff-complete. (C and D), We assign player 1 a fixed strategy from a space that is neither best-reply-complete nor payoff-complete, the space of reactive-2 strategies. This strategy is $\mathbf{p} = (1, 0.85, 0.85, 0)$, corresponding to the last two actions of the opponent, CC, CD, DC, DD. (C), Player 2 samples a large number of reactive-2 strategies, generating an area of feasible payoffs. Note that \mathbf{p} is the best-reply to itself in the space of reactive-2 strategies. (D), Among the self-reactive-2 strategies, one can find a better reply, such as the Alternator strategy that alternates between cooperation and defection each round. The resulting payoff for the Alternator is 2.05, which exceeds the payoff 2.00 of \mathbf{p} against itself. The area of feasible payoffs is strictly larger, showing that the space of reactive-2 strategies is not payoff-complete. The area on the right cannot grow any larger by expanding the strategy space of player 2: one can always find a payoff-preserving reply among the self-reactive-2 strategies.

Discussion

Repeated games and stochastic games play an important role in evolutionary game theory. They serve as natural models to capture the essence of reciprocity (6). Such models can explain the dynamics of friendships (40), competing firms (41), or of trade networks (13). In practice, individuals in these applications might employ highly nontrivial strategies, which depend on events arbitrarily far in the past. Yet when describing such interactions mathematically, researchers often restrict the strategies' complexity. Such a restriction can be useful. It may allow researchers to obtain analytical results that otherwise would be infeasible. At the same time, such restrictions also entail a risk. Results thus obtained might be spurious. Certain strategies might only evolve because mutants that could easily invade are rendered impossible by design. To address that risk, we introduce two complementary concepts of complete strategy spaces. If a space \mathbf{S} is best-reply-complete, then any strategy in \mathbf{S} has a best reply in \mathbf{S} . If the space is payoff-complete and the focal player uses a strategy in \mathbf{S} , then one can find for any strategy of the coplayer a strategy in \mathbf{S} that leaves the payoffs for both players unchanged. That is, any outcome that can be realized against an \mathbf{S} opponent can be realized with a strategy in \mathbf{S} . Levinský et al. (23) have introduced a general

criterion for a strategy space to be best-reply-complete. We show that the same criterion also implies payoff-completeness. Interestingly, when that criterion applies, the strategy space is payoff-complete irrespective of the precise game being considered and irrespective of the precise discount rate.

Complete spaces are useful when interpreting evolutionary simulations and determining best responses. For example, if a strategy is a best reply within a best-reply-complete strategy space, then it remains a best reply even if more complex strategies are considered.

Our motivation to study complete strategy spaces is related to recent arguments that strategy spaces should be unbiased (42). This latter requirement means that researchers should not cherry-pick the strategies they choose to compare. For example, instead of analyzing the competition of ALLC, ALLD, and Tit-for-Tat, researchers should permit all strategies of a given complexity class (in this case, say all reactive strategies). Our results, however, suggest that this requirement of unbiasedness may not be enough. Even when exploring the evolutionary dynamics among all reactive strategies, the winning strategy may still perish once we allow for strategies of higher or different complexity (Fig. 5B). To manage that risk, our two notions of completeness seem key. For example, when a given resident strategy $s \in \mathbf{S}$ is challenged by rare

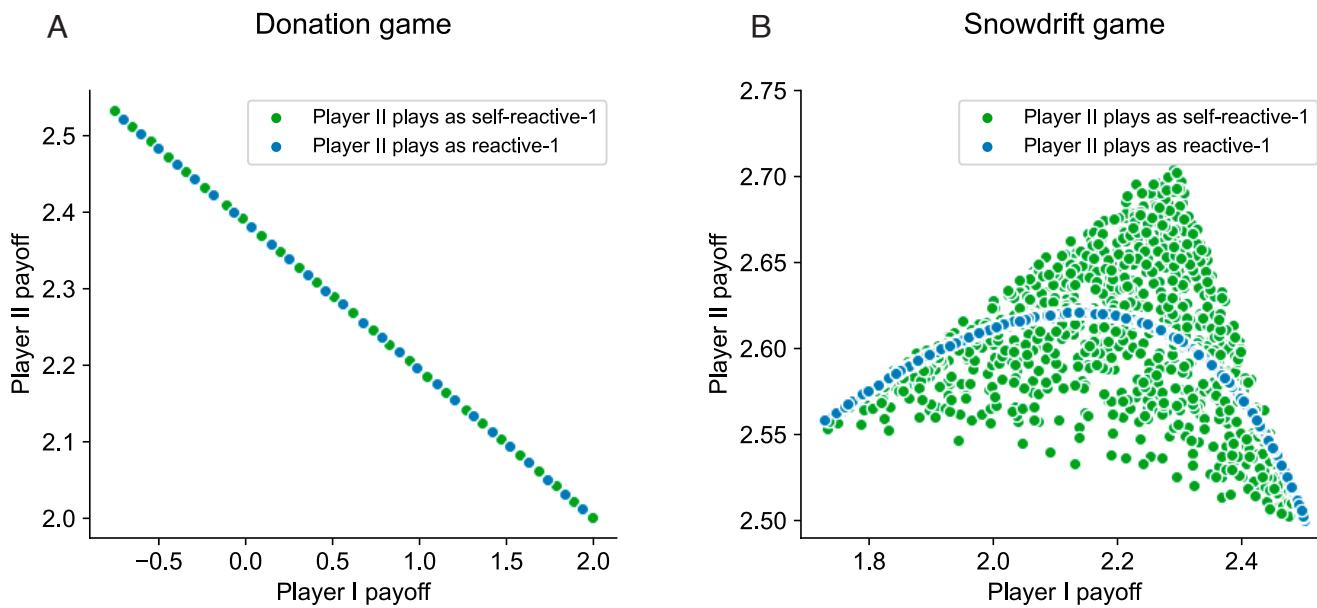


Fig. 5. Whether a space is complete can depend on the game. For spaces that are not based on recursive symmetric partitions, it may depend on the specific game whether the space is payoff-complete. Here, we assume player 1 adopts a fixed reactive-1 strategy $\mathbf{p} = (1, 0.85)$ (for a coplayer who cooperated or defected in the previous round, respectively). The partition associated to reactive-1 strategies is not symmetric. For the coplayer, we either sample reactive-1 strategies (blue points), or self-reactive-1 strategies (green points). (A), When payoffs are given by a donation game, $(a, b, c, d) = (2, -1, 3, 0)$, the space of reactive-1 strategies is payoff-complete. (B), For a snowdrift game, $(a, b, c, d) = (2.5, 2, 3, 0)$, the space is no longer payoff-complete. In both cases, for convenience, we have taken a discount rate approaching one, $\beta \rightarrow 1$.

mutants, payoff completeness of \mathbf{S} implies that the resident faces the full range of possible mutant payoffs. Any payoff that can be realized in an interaction against the resident can be realized by a mutant strategy drawn from within \mathbf{S} .

Importantly, however, this result does not ensure that the distribution of mutant payoffs is preserved. For example, for a given memory-1 resident, the generated payoffs may differ depending on whether random mutants are drawn from the space of memory-1 strategies or memory-2 strategies (even though any memory-2 strategy can be emulated within the memory-1 space). In particular, even for a complete strategy space, simulation results still depend on the specific process used to generate random mutants. Therefore, it remains important to explore the effect of different mutation schemes when performing evolutionary simulations.

The criterion for completeness involves the notion of recursive factors (23) or partitions. All partitions of memory-1 histories are recursive. Only a small number of memory- k partitions are recursive. However, every partition can be refined into a recursive partition. Therefore, for every strategy space, it is possible to find a (slightly) larger one which is complete. In *SI Appendix*, the general mathematical theory, which may be useful in further extensions, is discussed at length. We describe how recursiveness enters the picture when studying games, and what can go wrong when it is absent. We also describe an explicit algorithm for refining a partition to make it recursive. Many additional examples and details are also contained there.

An interesting open problem is to give a set of necessary and sufficient conditions for completeness. The condition we present here is sufficient but not necessary. Another open problem is whether certain strategies such as ALLD are necessarily part of any complete space. A last promising direction is to extend the idea of completeness to repeated games with an alternating move structure (43–45), or to extend payoff-completeness to other settings where some results on best replies are known, for

example, continuous strategy spaces (23, 46, 47), or spaces of time-dependent strategies (23, 48).

Materials and Methods

To establish our result, we use the theory of Markov decision processes, or MDPs. A Markov decision process is equivalent to a stochastic game with only one player. Beginning with a stochastic game, we expand the state space into a space of game histories. Thus, we track previous moves as well as the current state. This allows us to abstract out one player by absorbing this player's strategy into the transition probability. The result is a Markov decision process for the remaining player. We generate a multiobjective Markov decision process (or MOMDP) by storing both players' payoffs in a vector.

This technique is described at length in our *SI Appendix*. We derive our main result by applying a lemma which holds for a certain class of state-aggregated multiobjective Markov decision processes, which we call factored (MO)MDPs. A factored strategy for the abstracted player leads to a factored MDP for the focal player. The lemma states that any strategy for a factored MDP has a stationary (based on the current state) strategy which achieves the same vector payoff.

Data, Materials, and Software Availability. There are no data underlying this work.

ACKNOWLEDGMENTS. This material is based on work supported by the NSF Graduate Research Fellowship Program under Grant No. DGE 2146752. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF. C.H. and N.E.G. acknowledge the generous support by the European Research Council Starting Grant 850529: E-DIRECT.

Author affiliations: ^aDepartment of Mathematics, University of California, Berkeley, CA 94720; ^bDepartment of Mathematics, Harvard University, Cambridge, MA 02138; ^cMax Planck Research Group 'Dynamics of Social Behavior', Max Planck Institute for Evolutionary Biology, Plön 24306, Germany; ^dResearch Group on the Dynamics of Social Behavior, IT:U Interdisciplinary Transformation University, Linz A-4040, Austria; ^eDiscrete Event Simulation Research Team, RIKEN Center for Computational Science, Kobe 650-0047, Japan; and ^fDepartment of Organismic and Evolutionary Biology, Harvard University, Cambridge, MA 02138

1. R. L. Trivers, The evolution of reciprocal altruism. *Q. R. Biol.* **46**, 35–57 (1971).
2. G. J. Mailath, L. Samuelson, *Repeated Games and Reputations* (Oxford Univ. Press, Oxford, UK, 2006).
3. A. Rapoport, A. M. Chammah, *Prisoner's Dilemma* (University of Michigan Press, Ann Arbor, 1965).
4. J. Hofbauer, K. Sigmund, *Evolutionary Games and Population Dynamics* (Cambridge University Press, Cambridge, UK, 1998).
5. M. A. Nowak, Five rules for the evolution of cooperation. *Science* **314**, 1560–1563 (2006).
6. K. Sigmund, *The Calculus of Selfishness* (Princeton Univ. Press, Princeton, NJ, 2010).
7. P. Dal-Bó, G. R. Fréchette, On the determinants of cooperation in infinitely repeated games: A survey. *J. Econ. Lit.* **56**, 60–114 (2018).
8. J. García, M. van Velen, No strategy can win in the repeated prisoner's dilemma: Linking game theory and computer simulations. *Front. Robot. AI* **5**, 102 (2018).
9. C. Hilbe, K. Chatterjee, M. A. Nowak, Partners and rivals in direct reciprocity. *Nat. Hum. Behav.* **2**, 469–477 (2018).
10. N. Glynatsi, V. Knight, A bibliometric study of research topics, collaboration and centrality in the field of the Iterated Prisoner's Dilemma. *Humanit. Soc. Sci. Commun.* **8**, 45 (2021).
11. M. K. Schweinfurth, J. Call, Reciprocity: Different behavioural strategies, cognitive mechanisms and psychological processes. *Learn. Behav.* **47**, 284–301 (2019).
12. C. Rossetti, C. Hilbe, Direct reciprocity among humans. *Ethology* **130**, e13407 (2023).
13. M. R. Frank *et al.*, Detecting reciprocity at a global scale. *Sci. Adv.* **4**, eaao5348 (2018).
14. L. S. Shapley, Stochastic games. *Proc. Natl. Acad. Sci. U.S.A.* **39**, 1095–1100 (1953).
15. M. A. Nowak, K. Sigmund, Tit for tat in heterogeneous populations. *Nature* **355**, 250–253 (1992).
16. L. A. Imhof, M. A. Nowak, Stochastic evolutionary dynamics of direct reciprocity. *Proc. R. Soc. B Biol. Sci.* **277**, 463–468 (2010).
17. M. A. Nowak, K. Sigmund, A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner's Dilemma game. *Nature* **364**, 56–58 (1993).
18. C. Hauert, H. G. Schuster, Effects of increasing the number of players and memory size in the iterated prisoner's dilemma: A numerical approach. *Proc. R. Soc. B* **264**, 513–519 (1997).
19. F. L. Pinheiro, V. V. Vasconcelos, F. C. Santos, J. M. Pacheco, Evolution of all-or-none strategies in repeated public goods dilemmas. *PLoS Comput. Biol.* **10**, 1–5 (2014).
20. A. J. Stewart, J. B. Plotkin, Small groups and long memories promote cooperation. *Sci. Rep.* **6**, 26889 (2016).
21. C. Hilbe, L. A. Martinez-Vaquero, K. Chatterjee, M. A. Nowak, Memory- n strategies of direct reciprocity. *Proc. Natl. Acad. Sci. U.S.A.* **114**, 4715–4720 (2017).
22. Y. Murase, S. K. Baek, Grouping promotes both partnership and rivalry with long memory in direct reciprocity. *PLoS Comput. Biol.* **19**, e1011228 (2023).
23. R. Levinsky, A. Neyman, M. Zelený, Should I remember more than you? Best responses to factored strategies. *Int. J. Game Theory* **49**, 1105–1124 (2020).
24. 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).
25. A. J. Stewart, J. B. Plotkin, From extortion to generosity, evolution in the iterated prisoner's dilemma. *Proc. Natl. Acad. Sci. U.S.A.* **110**, 15348–15353 (2013).
26. A. Szolnoki, M. Perc, Defection and extortion as unexpected catalysts of unconditional cooperation in structured populations. *Sci. Rep.* **4**, 5496 (2014).
27. A. McAvoy, C. Hauert, Autocratic strategies for iterated games with arbitrary action spaces. *Proc. Natl. Acad. Sci. U.S.A.* **113**, 14934–14939 (2016).
28. M. Ueda, Memory-two zero-determinant strategies in repeated games. *R. Soc. Open Sci.* **8**, 202186 (2020).
29. X. Chen, F. Fu, Outlearning extortioners: Unbending strategies can foster reciprocal fairness and cooperation. *PNAS Nexus* **2**, pgad176 (2023).
30. A. McAvoy *et al.*, Unilateral incentive alignment in two-agent stochastic games. *Proc. Natl. Acad. Sci. U.S.A.* **122**, e2319927121 (2025).
31. M. Doebeli, C. Hauert, Models of cooperation based on the prisoner's dilemma and the snowdrift game. *Ecol. Lett.* **8**, 748–766 (2005).
32. F. Fu, M. A. Nowak, C. Hauert, Invasion and expansion of cooperators in lattice populations prisoner's dilemma vs. snowdrift games. *J. Theor. Biol.* **266**, 358–366 (2010).
33. C. Hilbe, A. Traulsen, K. Sigmund, Partners or rivals? Strategies for the iterated prisoner's dilemma. *Games Econ. Behav.* **92**, 41–52 (2015).
34. E. Akin, "The iterated prisoner's dilemma: Good strategies and their dynamics" in *Ergodic Theory, Advances in Dynamics*, I. Assani, Ed. (de Gruyter, Berlin, 2016), pp. 77–107.
35. D. P. Kraines, V. Y. Kraines, Learning to cooperate with pavlov an adaptive strategy for the iterated prisoner's dilemma with noise. *Theory Decis.* **35**, 107–150 (1993).
36. P. Molander, The optimal level of generosity in a selfish, uncertain environment. *J. Confl. Resolut.* **29**, 611–618 (1985).
37. N. E. Glynatsi, E. Akin, M. A. Nowak, C. Hilbe, Conditional cooperation with longer memory. *Proc. Natl. Acad. Sci. U.S.A.* **121**, e2420125121 (2024).
38. P. LaPorte, C. Hilbe, M. A. Nowak, Adaptive dynamics of memory-one strategies in the repeated donation game. *PLoS Comput. Biol.* **19**, 1–31 (2023).
39. F. Lesigang, C. Hilbe, N. Glynatsi, Can I afford to remember less than you? Best responses in repeated additive games. *Econ. Lett.* **250**, 112300 (2025).
40. J. Silk, "Cooperation without counting - The puzzle of friendship" in *Genetic and Cultural Evolution of Cooperation*, P. Hammerstein, Ed. (MIT Press, 2003), pp. 37–54.
41. D. Bernheim, M. D. Whinston, Multimarket contact and collusive behavior. *RAND J. Econ.* **21**, 1–26 (1990).
42. J. García, A. Traulsen, Picking strategies in games of cooperation. *Proc. Natl. Acad. Sci. U.S.A.* **122**, e2319925121 (2025).
43. M. A. Nowak, K. Sigmund, The alternating prisoner's dilemma. *J. Theor. Biol.* **168**, 219–226 (1994).
44. M. R. Frean, The prisoner's dilemma without synchrony. *Proc. R. Soc. B* **257**, 75–79 (1994).
45. P. S. Park, M. A. Nowak, C. Hilbe, Cooperation in alternating interactions with memory constraints. *Nat. Commun.* **13**, 737 (2022).
46. T. Killingback, M. Doebeli, N. Knowlton, Variable investment, the continuous prisoner's dilemma, and the origin of cooperation. *Proc. R. Soc. B* **266**, 1723–1728 (1999).
47. T. Killingback, M. Doebeli, The continuous prisoner's dilemma and the evolution of cooperation through reciprocal altruism with variable investment. *Am. Nat.* **160**, 421–438 (2002).
48. V. Huebner, C. Hilbe, M. Staab, M. Kleshnina, K. Chatterjee, Time-dependent strategies in repeated asymmetric public goods games. *Dyn. Games Appl.* **15**, 1617–1645 (2025).