# Evolutionary parsimony: an equilibrium refinement that sharply constrains the space of outcomes in games with multiple equilibria

Jean-Baptiste André[1]

[1]Institut Jean Nicod, Département d'études cognitives, ENS, EHESS, PSL Research University, CNRS, Paris France

## Abstract

Evolutionary game theory loses much of its predictive power in games with multiple equilibria. For such games, this paper introduces a simple and general refinement principle, grounded in evolutionary dynamics, that sharply narrows the set of possible outcomes. Rather than designing strategies from scratch, evolution shapes them gradually through the accumulation of adaptive mutations, the vast majority of which have small effects. This process can be approximated heuristically by assuming that smaller-effect mutations always occur first, ignoring the unlikely possibility that larger-effect mutations appear earlier. This approximation leads to a principle of adaptive parsimony: at each step, evolution always follows the simplest possible path. As a result, most theoretically possible equilibria are actually unreachable, as they would require a transition where a large-effect mutation fixes despite a simpler alternative being available. What remains is a small subset of equilibria that seem intuitively reasonable from a biological perspective: those that (i) preserve ecological symmetry, (ii) do not rely on non-credible threats, and (iii) avoid the bizarre behavioral patterns predicted by the folk theorem in repeated games.

# 1 Introduction

The success of evolutionary game theory, alongside behavioral ecology, lies in its ability to impose strict constraints on what evolution can and cannot produce. The stringent conditions required for a strategy to be evolutionarily stable enable precise insights into a wide range of behaviors, from sexual selection to parent-offspring conflict and kin altruism (1–3).

Yet, in certain areas of behavior, this predictive power is weaker, as multiple evolutionarily stable or neutrally stable strategies coexist, making it impossible to single out a clear evolutionary outcome. In such cases, evolutionary game theory loses the precision it offers in other contexts, predicting a broad range of possible equilibria without providing a straightforward way to choose between them.

This is particularly evident in reciprocal cooperation, where a wide range of strategies—some highly counterintuitive and involving intricate behavioral patterns—can be equilibria as long as they elicit favorable responses from others, a result central to the theory of repeated games and known as the ”folk theorem” (4–8).

And this problem is not unique to repeated games. In bargaining games, for instance, where players negotiate how to divide a resource, a single unique equilibrium is rare. Instead, multiple equilibria often arise, with each side's strategy co-adapted to the other (9–12).

To address this issue, game theorists have long worked to narrow down the set of plausible equilibria, developing a range of refinement concepts (13, 14). Within this broader effort, some researchers have specifically explored why certain equilibria might be favored over others in an evolutionary context, identifying two main mechanisms. The first is group selection, which is thought to favor socially efficient equilibria (15–18). The second involves selection pressure from rare mutants, captured by the second ESS condition or the concept of robustness to indirect invasion, whose effects depend on the specific model (19–24).

However, these mechanisms have two shortcomings. First, they rely on forces that are generally considered weak in evolutionary terms. While they may influence allelic frequencies under restrictive conditions, these forces are typically overshadowed by other selective pressures (25–27). Second, they act ex post, focusing on competition between equilibria after they have emerged, leaving unanswered the question of which equilibria are more likely to arise in the first place.

This paper takes a different approach. Rather than focusing on weak forces acting through ex post competition, it aims to understand how evolutionary dynamics narrow the range of equilibria ex ante, with some more likely to be reached than others.

While a rich body of literature explicitly models evolutionary dynamics in repeated games (28–31), it has yet to yield a simple, overarching theory. Instead, each case appears unique, with factors like the choice of mutation matrix playing a decisive role (32, 33).

Such an overarching theory already exists for the evolution of *quantitative* traits: the theory of adaptive dynamics (34–38). It provides broad insights into evolutionary processes based on two central assumptions: adaptive mutations are rare and have small effects. These assumptions enable two approximations that make evolutionary modeling tractable. First, the rarity of adaptive mutations allows them to be treated as occurring sequentially, with each mutation either fixing or being lost before the next one appears. This allows to approximate evolution as a 'trait substitution sequence.' Second, the small effect of these mutations allows the adaptation process to be approximated as continuous, with each mutation producing an infinitesimally small change.

This paper seeks to extend these assumptions to games with a discrete strategy space, such as repeated games. The challenge is that the assumption of small mutation effects cannot be directly applied, as it would require setting an arbitrary upper limit on mutation size. To address this, I propose a gradual interpretation of the assumption: instead of imposing a fixed limit on mutation size, it simply posits that while any mutation is possible, smaller-effect mutations occur more frequently than larger ones. This assumption is strongly supported by both theoretical models of adaptation (39–41) and empirical evidence on the distribution of beneficial mutation effects (42–44), both of which show that small-effect mutations are more frequent than large-effect ones.

This extension allows for a different type of approximation suited to discrete models. At each step in the trait substitution sequence, when multiple substitution events are theoretically possible—because several mutant strategies offer an advantage over the resident—it can be approximated that the mutation with the smallest effect will always arise first. Consequently, at each step, evolution consistently favors the transition that requires the fewest mutational changes. In other words, while the small mutation effects assumption for quantitative traits allows adaptation to be approximated as a gradual process, its extension to discrete traits allows adaptation to be approximated as proceeding in the 'most gradual way possible.'

This approximation narrows the range of possible evolutionary dynamics, providing heuristic insights into which equilibria are most likely to arise from a given initial state. Its main effect is to act as an evolutionary Occam's razor: when multiple evolutionary transitions are possible, the simplest one always occurs. Larger evolutionary changes can happen only if subsequent mutations provide an advantage after the simplest mutant has already established itself. For this reason, I refer to the dynamics and equilibria resulting from this approximation as evolutionarily parsimonious.

The multiplicity of equilibria arises from the social nature of selection: the payoff of an action depends on how others respond, and their payoff, in turn, depends on further reactions. In some games, this circularity weakens the connection between selection and environmental constraints, leaving room for arbitrariness in adaptive behavior. Any behavioral pattern, however complex, can become an equilibrium as long as it elicits favorable responses from others.

The main effect of evolutionary parsimony is to prevent the evolution of such complex patterns. Since they have no intrinsic advantage, their only benefit being to match similar patterns in others, evolutionary transitions, under the parsimony approximation, will never produce them. Instead, evolution will always favor simpler mutations that are equally advantageous but free of arbitrary complexity.

As a result, parsimonious equilibria make up only a tiny fraction of Nash equilibria, and they are defined by a simplicity that naturally aligns with what seems reasonable and intuitive to a biologist. (1) By default, parsimonious equilibria preserve ecological symmetry, meaning that individuals with identical action sets and payoff functions adopt the same behaviors at equilibrium. (2) Parsimonious equilibria rule out non-credible threats. And (3) parsimonious equilibria are free from the bizarre behavioral patterns that make up the vast diversity of equilibria in repeated games.

# 2 Evolutionarily parsimonious equilibria

Consider a game defined by a finite set of $n$ players, a tree representing the sequence of decisions, a set of terminal nodes with associated payoffs for each player, and a partition of the non-terminal nodes into $n+1$ subsets—one for each player and one for "Nature's moves", representing events beyond the players' control. A player's strategy is defined by a mapping from the set of their nodes to the set of available behavioral policies.

Following the standard assumption of vanishingly rare mutations, the evolution of strategies in this game is approximated as a trait substitution sequence. Each substitution involves a mutant emerging within a resident population, surviving initial stochastic phases, and ultimately becoming fixed. At each step, multiple substitutions are possible, as several mutations may have a non-zero probability of fixation. Evolution is thus a stochastic process with multiple possible realizations.

To simplify the analyses, two further assumptions are made. First, the mutation graph is assumed complete, meaning any strategy can potentially arise as a mutation from any other. Second, an infinite population is assumed, allowing only substitutions to strictly advantageous strategies, as mutants with neutral or deleterious effects have zero fixation probability.

A final assumption is made for biological realism. The effect of selection from rare mutants is neglected, meaning only mutants strictly favored over the resident can invade, while those neutral against the resident but favored when competing against other mutants cannot. The rationale is that mutants strictly neutral against the resident represent a degenerate case unlikely to occur in biological systems. Selection driven by rare mutants is therefore largely a theoretical artifact, often leading to unrealistic scenarios, particularly in repeated games. However, in the Supporting Information (SI) F, I show that the concept of parsimony remains applicable without this assumption, yielding nearly identical results.

Under these assumptions, the evolutionary transition from strategy $S_1$ to strategy $S_2$ is said to be *feasible* if and only if $P(S_2, S_1) > P(S_1, S_1)$, where $P(X, Y)$ denotes the payoff for an individual using strategy $X$ when interacting with an individual using $Y$. This means that $S_2$ must perform strictly better in interactions with $S_1$.

At each step of the stochastic trait substitution sequence, multiple transitions may be feasible, but most will not be parsimonious (see fig. 1 for a visual illustration). Formally, for each pair of strategies $S_1$ and $S_2$, a positive function $d(S_1, S_2)$ is defined to represent the mutational distance from $S_1$ to $S_2$: the greater this distance, the lower the probability that random biochemical events will transform a genotype coding for $S_1$ into one coding for $S_2$. In the applications discussed in this paper, mutational distances will be measured using finite automata (see SI A for details). However, any measure of distance can be used in principle.

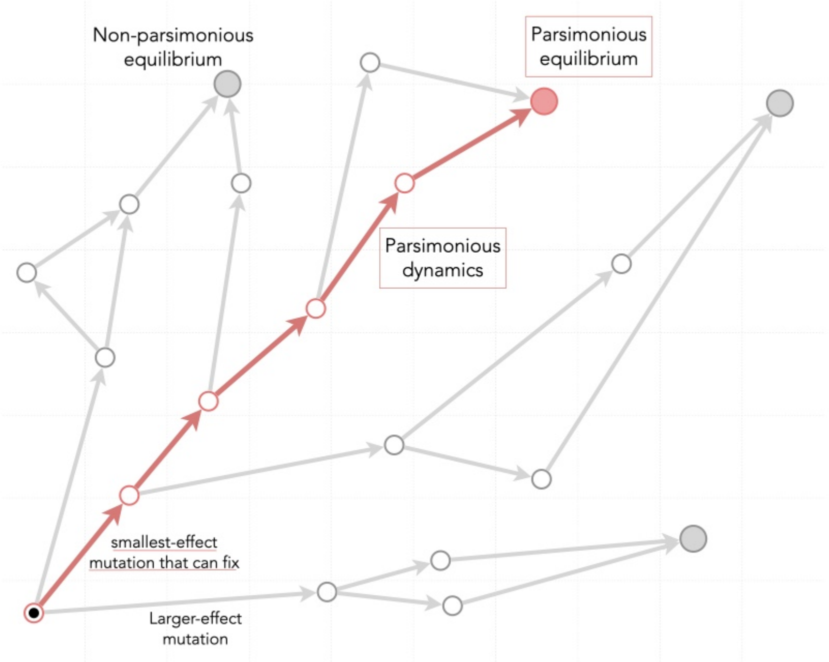

Figure 1: **Parsimonious vs non-parsimonious dynamics and resulting equilibria in strategy space.** Evolutionary trajectories are shown from an ancestral strategy (circle with central dot). Parsimonious path (red) and non-parsimonious paths (grey) lead to different equilibrium outcomes (filled circles).

The evolutionary transition from strategy $S_1$ to strategy $S_2$ is then said to be *parsimonious* iff (i) $P(S_2, S_1) > P(S_1, S_1)$ (i.e., the transition is feasible), and (ii) the following condition is met:

$$\forall S_3 \neq S_2 \text{ with } P(S_3, S_1) > P(S_1, S_1),\, d(S_3, S_1) \geq d(S_2, S_1) \tag{1}$$

In other words, $S_2$ is the closest strategy that can successfully invade $S_1$ (see fig. 1).

A strategy $S^*$ is then defined as an *evolutionarily parsimonious equilibrium* from an ancestral strategy $S_0$ if and only if the two following conditions hold:

1. There exists at least one trait substitution sequence from $S_0$ to $S^*$ that includes only parsimonious transitions.

2. $\forall S \neq S^* P(S, S^*) < P(S^*, S^*)$, i.e., $S^*$ is an equilibrium.

The parsimonious approximation assumes that evolution can reach only parsimonious equilibria and never non-parsimonious ones. While this is a potentially strong approximation, we will now see that it offers significant heuristic value in practice, in games with multiple equilibria.

# 3 Parsimony narrows the range of evolutionary equilibria

The constraints imposed by parsimony on evolutionary equilibria follow a simple logic. Here, I outline this logic in broad terms before turning to a formal analysis in three specific cases.

Arbitrary behavioral patterns, characteristic of games with many equilibria, always involve strategies that adjust behavior based on cues entirely unrelated to payoffs. In repeated games, for instance, this includes strategies that adjust their behavior depending on the round number.

To implement such strategies, individuals must be able to recognize and respond to these arbitrary cues, which is not a given. Doing so requires underlying biological mechanisms. As a result, evolving a response to arbitrary cues is likely to come with a 'cost' in terms of mutational distance. Starting from a hard-wired ancestral strategy that cannot respond to a given cue, the mutational step required to evolve a conditional response is larger than that needed to evolve another hard-wired strategy.

The parsimony approximation therefore imposes a constraint on the evolution of such conditional dispositions. They can only evolve if they provide a benefit unattainable through simpler strategies closer to the ancestral form.

Because arbitrary conditional dispositions never provide such a benefit, they are systematically excluded under parsimony. If a resident population lacks such a disposition and could, in theory, be invaded by a mutant strategy that possesses it, there is always a simpler mutant—without the arbitrary conditionality—that can invade instead. As a result, the set of equilibria that evolution can genuinely reach in practice is only a small fraction of the full range of theoretically possible equilibria (see fig. 2 for a schematic representation).

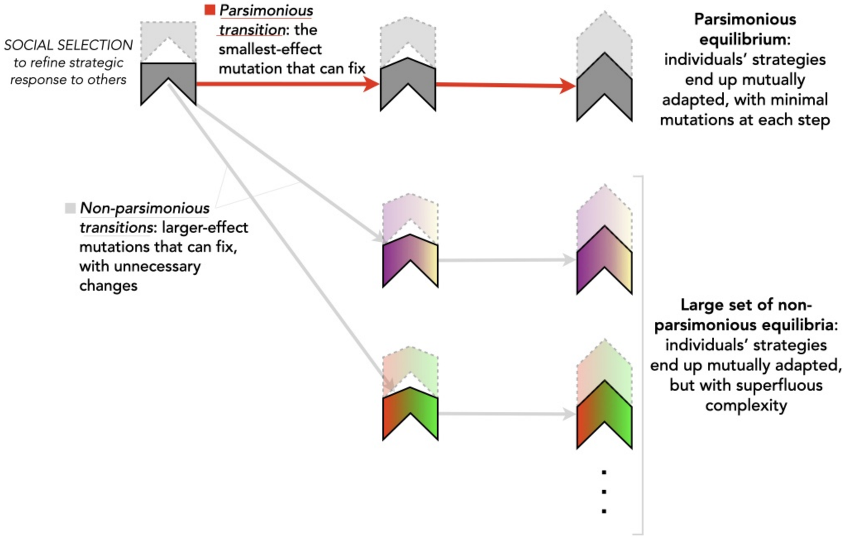

Figure 2: **Schematic representation of how parsimony narrows the range of evolutionary equilibria.** Starting from an ancestral strategy where individuals respond poorly to one another, selection favors refinements that improve strategic response to others. The smallest such change is shown here as a shift in shape. A wide range of other evolutionary transitions are also possible, where strategic coordination is achieved not only through the minimal change in shape, but also through superfluous conditional dispositions (represented by matching colors), but these alternatives are not parsimonious. As a result, the range of parsimonious equilibria is only a small fraction of the full range of theoretically possible equilibria.

To illustrate how this constraint shapes evolutionary outcomes, I now examine three specific consequences.

## 3.1 By default, evolution preserves ecological symmetry

In many games, the concepts of Nash equilibrium and Evolutionarily Stable Strategy permit a wide range of arbitrarily asymmetric equilibria. Even when individuals are fully symmetrical in ecologically relevant terms—sharing the same set of feasible actions and payoffs—they can use any perceptible asymmetry in their environment, even if it is entirely unrelated to payoffs, to condition their strategies and adopt asymmetric behaviors at equilibrium (45).

From a biological perspective, this is counterintuitive. It is hard to see why individuals would complicate their behavior by conditioning their actions on arbitrary symmetry-breaking cues merely because others do the same.

The principle of parsimony reflects this biological intuition. In any symmetric game with $n$ players, if the ancestral strategy is fully symmetric—meaning players do not condition their actions on arbitrary cues—then all parsimonious equilibria are symmetric.

The reasoning is simple (see SI B for details). An asymmetric strategy relies on the evolution of a conditional capacity, allowing behavior to vary based on an arbitrary environmental cue that serves as a symmetry breaker. Such asymmetry comes at a cost in terms of mutational distance. For an asymmetric strategy to evolve parsimoniously, the advantage it provides would need to be unattainable by any simple symmetric strategy. Yet, by definition, against a symmetric resident, behaving asymmetrically offers no inherent benefit. Any advantage an asymmetric mutant gains in one state can be achieved just as effectively by applying the same action across all states, since the symmetric resident treats all states identically. As a result, under the parsimony approximation, a symmetric mutant will always invade before an asymmetric one, preventing any evolutionary symmetry breaking.

Evolution can nevertheless produce asymmetries, but only under two conditions. First, there must be a slight initial ecological asymmetry—whether in payoffs, the action set, or the behavior of the ancestral strategy. Second, exagerating this initial asymmetry must provide an immediate advantage to individuals.

This can be formalized with a simple scenario where two individuals compete over an indivisible resource, modeled as a hawk-dove game (45; see SI B for details). Each individual chooses whether to compete for the resource or abstain entirely, with a cost incurred if both choose to compete. Now, suppose a slight initial ecological asymmetry exists—for instance, the cost of conflict is slightly lower for one player, with the difference linked to a perceptible environmental feature.

Consider an ancestral symmetric resident strategy that plays Hawk with a fixed probability $p \in [0, 1]$, regardless of the player's side. Under parsimonious dynamics, this symmetry is initially preserved, and the population converges toward a strategy where Hawk is played with a probability $\hat{p} \equiv \frac{T-R}{T+S-P-R}$ in both states. Once this threshold value is reached, however, all symmetric strategies become strictly neutral, leaving only asymmetric mutants with the potential to gain a selective advantage.

At this point, if one state reduces the cost of conflict, playing Hawk more frequently in that state becomes advantageous, driving selection to further reduce Hawk play in the other state. This amplification continues until it results in a fully asymmetric equilibrium where only one individual claims the resource. Hence, the final asymmetry far exceeds the initial asymmetry.

By contrast, in a different game where asymmetry is not needed for coordination (see SI C), selection does not amplify the initial perturbation, leaving the final equilibrium asymmetry no greater than the slight perturbation introduced initially.

This shows that quasi-conventional symmetry breaking can evolve, as long proposed in evolutionary game theory (46), but only under specific conditions. First, asymmetry cannot arise arbitrarily; there must be an initial difference in payoffs or strategies to break symmetry in the first place. Second, evolution will only amplify a small initial asymmetry if it provides coordination benefits that cannot be achieved through symmetry, such as reducing conflict, enabling division of labor, or supporting niche differentiation.

## 3.2 Evolution does not lead to strategies entailing non-credible threat

One of the most notorious flaws of the Nash equilibrium concept, and its evolutionary counterpart, the neutrally stable strategy (as the ESS concept does not apply in this situation), is their inability to exclude non-credible threats. The mere possibility of threatening others with punishment can allow an individual to extract benefits in equilibrium, even when it is clear they have no incentive to follow through on the threat (11, 47).

From a biological perspective, this is paradoxical. Evolution shapes mechanisms according to the benefits they provide. If an action offers no advantage and only incurs costs, the disposition to perform it should not be produced by selection, regardless of whether it might serve as a threat.

To address this apparent paradox, evolutionary game theorists typically turn to the concept of subgame perfection from standard game theory (13). A subgame perfect equilibrium is a strategy that is not only the best response to itself, as in a Nash equilibrium, but also the best response in every possible contingency, including those that should never arise in equilibrium, thereby ruling out non-credible threats.

The problem is that translating the concept of subgame perfection into an evolutionary framework is not straightforward. Biological evolution has no reason to shape strategies for hypothetical contingencies that never arise. Consequently, applying subgame perfection to evolutionary settings relies on assumptions that artificially introduce selection pressures along off-equilibrium paths, whether by assuming behavioral errors or invoking selection due to the presence of rare mutants (48–50).

The principle of parsimony offers an alternative solution to the same problem—one that is simpler, more biologically grounded, and independent of subgame perfection.

This can be formalized using the ultimatum game, in which one player, the proposer, offers a fraction of a resource to another player, the responder, who then chooses whether to accept or reject the offer (51; see SI D for details). The responder can adopt one of three types of strategies: two extreme, hard-wired strategies—accepting all offers or rejecting all offers regardless of their value—and a range of intermediate strategies that are conditional, accepting some offers while rejecting others.

Biologically speaking, intermediate strategies are distinct because they require a capacity to evaluate offers and adjust behavior based on their value. This makes it reasonable to assume that the mutational distance between the two extreme strategies is smaller than the distance from either extreme to an intermediate strategy. A mutation from one extreme to the other requires only a change in hard-wired behavior, while a mutation leading to an intermediate strategy requires two changes: the emergence of a new behavior and the additional capacity for conditionality.

The capacity for conditionality, however, provides no inherent advantage. A mutant that conditionally accepts some offers and rejects others may be favored in certain situations, but its advantage is always driven by its ability to accept some offers—not by its capacity to reject others. Thus, whenever a conditional acceptor is favored, a simpler mutant that unconditionally accepts all offers, regardless of their value, will always be at least as favored, if not more so. Under the parsimony approximation, this simpler mutant will always invade first, blocking any evolutionary transition toward conditional acceptance.

As a result, the only parsimonious equilibrium is one in which the responder accepts every offer, no matter how small. Parsimony eliminates the possibility of non-credible threats.

### 3.3 Evolution does not generate an extravagant diversity of equilibria in repeated games

In repeated games, the concepts of Nash equilibrium and subgame perfect equilibrium notoriously allow for a vast array of arbitrarily complex strategies (52). These include individuals cooperating only on specific rounds based on arbitrary patterns, conditioning their cooperation on environmental variables unrelated to the game's payoffs, or engaging in elaborate signaling systems and pre-cooperative behavioral sequences.

From a biological perspective, the possibility of individuals relying on such bizarrely complex equilibrium strategies is highly counterintuitive, appearing more like a theoretical artifact than a meaningful prediction. Once again, this intuition is captured by the principle of parsimony.

This is formalized here in a game referred to as the ”investment game,” introduced in a previous work (53) and designed to capture the ecological conditions underlying the evolution of reciprocal cooperation (54–57 ; see SI E for details; see also SI F.3 for an analysis of the standard repeated prisoner's dilemma, which yields identical results).

The investment game is an asymmetric repeated interaction between two individuals (see fig. 3). One decides whether to invest in their partner, while the other chooses whether to reciprocate at a personal cost. Before the first round, the investor faces uncertainty about the nature of the game. In some cases, investing yields an incidental benefit without requiring the partner to reciprocate—what we refer to here as a byproduct cooperation game (58–61). In others, the benefit arises only if the partner actively reciprocates; in other words, each round is a trust game. Finally, there are cases where the partner cannot offer any mutually beneficial reward at all, a situation we refer to as innefficient cooperation.

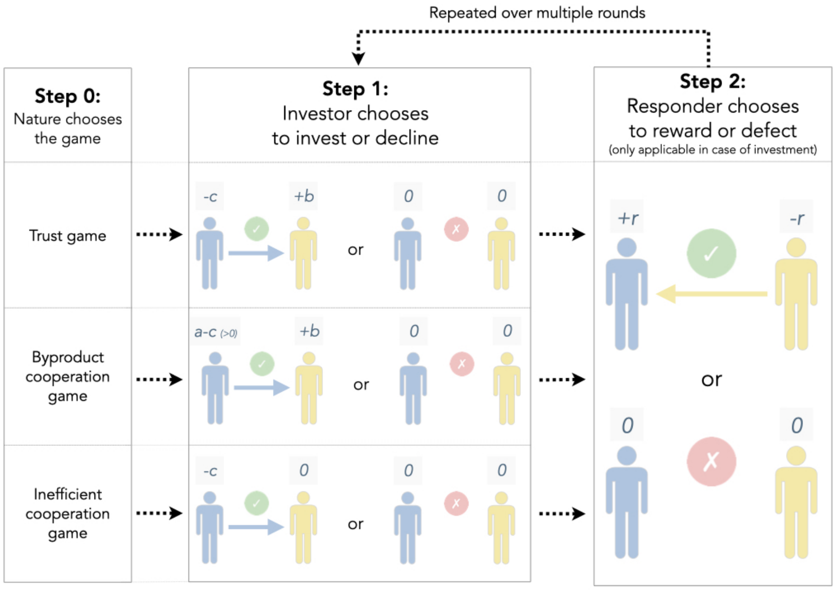

Figure 3: **Schematic representation of the investment game.**

This represents a scenario where conditional investment in cooperation can be advantageous, even in the absence of pre-existing reciprocal mechanisms, effectively bootstrapping the evolution of reciprocity (53, 54). Investors are initially selected to cooperate at least once to determine the type of game and to continue cooperating only if it proves beneficial. This conditional cooperation, in turn, creates selection pressure on their partners to actively reciprocate when the game is a trust game, ensuring the investor's continued cooperation in the future.

Like all repeated games, the investment game allows for a wide range of Nash equilibria. In all cases, players cooperate fully in byproduct games and never cooperate when cooperation would be wasteful. However, a broad spectrum of outcomes arises in trust games. At one extreme, reciprocity is entirely absent, and cooperation never occurs in trust games. At the other, reciprocal rewards are consistently provided, resulting in full cooperation in trust games. Between these extremes lies a vast array of intermediate strategies, where reciprocal rewards are given only in certain rounds of trust games, following arbitrary patterns, with investors adjusting their investments accordingly.

Yet, starting from a non-cooperative ancestral state, the constraint of parsimony eliminates all intermediate strategies, leaving only the two extreme equilibria as possible outcomes. First, intermediate strategies require the ability to condition behavior not only on receiving a reward but also on the round number. This makes them mutationally further from the non-cooperative ancestor than strategies that condition cooperation solely on receiving a reward. Second, while intermediate strategies can invade a non-cooperative resident, a simpler strategy—one that ignores the round number—can always invade just as effectively. Under the parsimony approximation, this simpler strategy, being closer in mutational distance to the ancestral state, will always invade first, blocking the evolution of intermediate strategies.

Thus, if reciprocal cooperation does evolve, it will not rely on arbitrary conditionality. Individuals will adjust their cooperation based on cues that carry real meaning—those that provide genuine information about the payoffs of cooperation. They will cooperate when it leads to a mutually beneficial outcome and refrain from cooperating when it would be ineffective. Contrary to the predictions of the folk theorem, they will never follow conditional rules that tie cooperation to meaningless cues with no connection to mutual benefits.

That said, reciprocal cooperation may not always evolve—selection for cooperation in byproduct games does not necessarily bootstrap cooperation in trust games (53). Whether it does depends on finer factors related to the availability of mutations. Unlike the hypothesis that arbitrarily complex conditional strategies are less likely to appear than simpler ones, which follows from well-grounded principles of parsimony, these considerations are far more tenuous, making it impossible to formulate general hypotheses. As a result, whether a particular form of reciprocal cooperation evolves through the bootstrapping of a particular form of byproduct cooperation will depend on the specific details of each case (53).

# 4 Discussion

The wide diversity of equilibria found in many games, especially repeated games, which undermines the predictive power of evolutionary game theory (4–8), stems from a common feature to all these games: the payoff of any given action depends largely on how others respond to it. As a result, almost any arbitrarily complex behavioral pattern can become an equilibrium, provided it elicits favorable responses from others.

In this article, I have argued that, within this extravagant diversity, only a small subset of equilibria can actually be reached through the process of biological evolution, and that this subset can be identified using a simple criterion, which I have formalized and called evolutionary parsimony.

An engineer can design a decision-making device from scratch, no matter how complex, and is therefore genuinely confronted with the full diversity of possible equilibria. Evolution, however, does not build decision-making device from scratch. Instead, it constructs them gradually, through the accumulation of small-effect mutations, each advantageous at the time it arises (62).

Such a gradual process can create remarkably complex and sophisticated traits, but only when they offer individuals a direct advantage in interacting with their environment. It cannot create arbitrarily complex traits whose only benefit comes from matching equally arbitrary complexities in others.

As a result, evolution can reach only a tiny fraction of the vast range of possible equilibria—those free of arbitrarily complex patterns, and therefore precisely the ones that seem reasonable and intuitive to a biologist. In repeated cooperation, for instance, these are the strategies where individuals cooperate when it is mutually beneficial and ignore arbitrary behaviors unrelated to payoffs.

While much of the literature has focused on how evolution might eliminate, ex post, the bizarre equilibria predicted by game theory—relying on weak forces like group selection (15–18) or selection by rare mutants (19–24), evolutionary parsimony suggests a simpler point: these equilibria never arise in the first place.

This analysis relies on an approximation of the evolutionary process. It assumes that when multiple evolutionary transitions are possible, the most likely one always occurs. In reality, however, the most likely transition is just that—more likely, but not certain. Strictly speaking, evolutionary parsimony should therefore be a quantitative concept rather than the all-or-nothing principle I have presented here.

Even so, this approximation is useful. It makes parsimony easy to apply, which is what gives it its practical heuristic value. Most importantly, it works well in practice because it is designed to rule out profoundly unreasonable equilibria, characterized by a qualitatively higher level of complexity, whose emergence through mutation is not just slightly less likely but extremely improbable. Put simply, it tells us: *a complex strategy is unlikely to evolve as a byproduct of selection for a simple one.* And that alone is enough to clarify which equilibria evolution is likely to reach.

I now conclude with a discussion of two standard equilibrium refinements from game theory, which share some similarities with the concept of parsimony but also have fundamental differences.

First, a class of game-theoretic approaches refines equilibria by taking into account the cost of complexity (20, 63–67). Similar to parsimony, the idea is that a strategy's complexity—particularly its conditional dispositions—must be outweighed by some benefits for the strategy to be an equilibrium.

The key difference is that this refinement only eliminates traits that are superfluous *in equilibrium*. In contrast, evolutionary parsimony constrains the *transitions* leading to complex traits. Even if a complex behavioral disposition—such as the ability to respond conditionally to a cue—could become strictly necessary once established as an equilibrium, its emergence would still require a non-parsimonious transition at some stage, making that equilibrium unreachable. This is why only parsimony can so drastically limit the diversity of equilibria in repeated games.

Second, another refinement from game theory, called subgame perfection, refines equilibria by requiring a strategy to be optimal in every possible contingency, even those that never occur in equilibrium (13). Like parsimony, subgame perfection is concerned with some form of optimality beyond the equilibrium path. As a result, in cases such as non-credible threats, both refinements yield the same outcome. From the perspective of subgame perfection, carrying out a non-credible threat is suboptimal. From the perspective of parsimony, evolution cannot even shape the ability to make such threats in the first place.

But the concepts of parsimony and perfection are also profoundly different. Perfection implies that a strategy responds optimally in all possible histories, assuming that the continuation game is played with a partner who plays the equilibrium strategy itself. In contrast, parsimony implies that the mechanisms involved in a strategy gave an advantage in interactions with individuals playing *other* strategies encountered on the path to equilibrium. This has two consequences.

First, parsimonious equilibria are often imperfect, since parsimony does not imply that a strategy responds optimally to *all* possible game histories. This makes sense for a concept that seeks to reflect the consequences of biological evolution. Any action taken in a history that has never occurred over the course of evolution is neutral and has no reason to be optimized by selection. That is, the concept of parsimony recognizes that evolution can lead to mismatches when organisms are placed in non-ecological situations.

Second, in repeated games, parsimony is even incompatible with perfection (see SI **??**). To achieve subgame perfection in repeated games, game theorists have introduced strategies known as Grim strategies, which punish themselves if they deviate (52). Even if one can appreciate the theoretical purpose of this trick, it is hard for a biologist not to find it absurd. And once again, this intuition is well captured by the principle of parsimony. While self-punishing strategies can be equilibria, the ability to self-punish never provides a selective advantage to a mutant in any situation. As a result, such mechanisms cannot evolve parsimoniously from an ancestral strategy that lacks them.

**Acknowledgments.** This study was supported by the EUR FrontCog grant ANR-17-EURE-0017 and the “Soutien à la mobilité internationale” of InSHS-CNRS. During the preparation of this work the author used ChatGPT4o in order to improve clarity and readability. After using this tool, the author reviewed and edited the content as needed and takes full responsibility for the content of the publication.

# References

[1] Jean-Baptiste Grodwohl and Geoff A. Parker. The early rise and spread of evolutionary game theory: perspectives based on recollections of early workers. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 378(1876):20210493, May 2023. ISSN 0962-8436, 1471-2970. doi: 10.1098/rstb.2021.0493. URL `https://royalsocietypublishing.org/doi/10.1098/rstb.2021.0493`.

[2] Olof Leimar and John M. McNamara. Game theory in biology: 50 years and onwards. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 378 (1876):20210509, May 2023. ISSN 0962-8436, 1471-2970. doi: 10.1098/rstb.2021.0509. URL `https://royalsocietypublishing.org/doi/10.1098/rstb.2021.0509`.

[3] Nicholas B. Davies, John R. Krebs, and Stuart A. West. *An Introduction to Behavioural Ecology*. John Wiley & Sons, 2012. ISBN 978-85-7811-079-6. doi: 10.1017/CBO9781107415324.004. arXiv: 1011.1669v3 ISSN: 1098-6596.

[4] Long-term Competition and Game-theoretic Analysis. Long-Term Competition—A Game-Theoretic Analysis with L. S. Shapley. pages 1–15, 1994.

[5] Fudenberg and Maskin 1986 The Folk Theorem in repeated games (Econometrica).pdf.

[6] D Fudenberg and J Tirole. *Game theory*. The MIT Press, Cambridge, MA, 1991.

[7] Boyd. Reciprocity: You have to think different. *Journal of Evolutionary Biology*, 19 (5):1380–1382, 2006. ISSN 1010061X. doi: 10.1111/j.1420-9101.2006.01159.x. ISBN: 1010061X.

[8] R Boyd and P J Richerson. Punishment Allows the Evolution of Cooperation (or Anything Else) in Sizable Groups. *Ethology and Sociobiology*, 13(3):171–195, 1992.

[9] John F Nash. The Bargaining Problem. 1950. doi: 10.4324/9780429303005-7.

[10] Ariel Rubinstein. Perfect Equilibrium in a Bargaining Model. *Econometrica*, 50(1):97, January 1982. ISSN 00129682. doi: 10.2307/1912531. URL https://www.jstor.org/stable/1912531?origin=crossref.

[11] John Gale, Kenneth G. Binmore, and Larry Samuelson. Learning to be imperfect: The ultimatum game. *Games and Economic Behavior*, 8(1):56–90, January 1995. ISSN 08998256. doi: 10.1016/S0899-8256(05)80017-X. URL https://linkinghub.elsevier.com/retrieve/pii/S089982560580017X.

[12] K. Binmore. Bargaining in biology? *Journal of Evolutionary Biology*, 23(7):1351–1363, 2010. ISSN 1010061X. doi: 10.1111/j.1420-9101.2010.02011.x. ISBN: 1010-061X.

[13] R Selten. Spieltheoretische Behandlung eines Otigopolmodells mit Nachfragetrfigheit. *Zeitschrift f{ü}r Gesamte Staatswissenschaft*, 121:301–324, 1965.

[14] J Harsanyi and R Selten. *A general theory of equilibrium selection in games*. MIT Press, 1988.

[15] Boyd and Richerson. Group Selection among Alternative Evolutionarily Stable Strategies. *Journal of Theoretical Biology*, 145(3):331–342, 1990.

[16] Boyd and Richerson. Group beneficial norms can spread rapidly in a structured population. *J Theor Biol*, 215(3):287–296, 2002. ISSN 0022-5193. doi: 10.1006/jtbi.2001.2515. ISBN: 0022-5193.

[17] Boyd and Richerson. Culture and the evolution of human cooperation. *Philosophical Transactions of the Royal Society B-Biological Sciences*, 364(1533): 3281–3288, November 2009. ISSN 1471-2970. doi: 10.1098/rstb.2009.0134. URL http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2781880&tool=pmcentrez&rendertype=abstract.

[18] Boyd and Richerson. Voting with your feet: Payoff biased migration and the evolution of group beneficial behavior. *Journal of Theoretical Biology*, 257(2):331–339, 2009. ISSN 00225193. doi: 10.1016/j.jtbi.2008.12.007. ISBN: 0022-5193.

[19] Drew Fudenberg and Eric Maskin. Evolution and Cooperation in Noisy Repeated Games. *American Economic Review*, 80(2):274 – 279, 1990. ISSN 00028282. URL http://www.jstor.org/stable/2006583. ISBN: 00028282.

[20] Binmore and Samuelson. Evolutionary stability in repeated games played by finite automata. *Journal of Economic Theory*, 57(2):278–305, August 1992. ISSN 00220531. doi: 10.1016/0022-0531(92)90037-I. URL http://linkinghub.elsevier.com/retrieve/pii/002205319290037I.

[21] M A Nowak, A Sasaki, C Taylor, and D Fudenberg. Emergence of cooperation and evolutionary stability in finite populations. *Nature*, 428(6983):646–650, 2004.

[22] André and Troy Day. Perfect reciprocity is the only evolutionarily stable strategy in the continuous iterated prisoner's dilemma. *Journal of Theoretical Biology*, 247 (1):11–22, July 2007. ISSN 0022-5193. doi: 10.1016/j.jtbi.2007.02.007. URL http://www.ncbi.nlm.nih.gov/pubmed/17397874.

[23] Matthijs Van Veelen. Robustness against indirect invasions. *Games and Economic Behavior*, 74(1):382–393, 2012. ISSN 08998256. doi: 10.1016/j.geb.2011.05.010. URL http://dx.doi.org/10.1016/j.geb.2011.05.010. Publisher: Elsevier Inc.

[24] Mathias Spichtig and Martijn Egas. When and How Does Mutation-Generated Variation Promote the Evolution of Cooperation? *Games*, 10(1):4, January 2019. ISSN 2073-4336. doi: 10.3390/g10010004. URL https://www.mdpi.com/2073-4336/10/1/4.

[25] George Christopher Williams. *Adaptation and Natural Selection: A Critique of Some Current Evolutionary Thought*. Princeton Science Library, 1966. URL https://www.amazon.com/Adaptation-Natural-Selection-Evolutionary-Princeton/dp/0691182868/ref=sr_1_1?dchild=1&keywords=Adaptation+and+Natural+Selection+williams&qid=1605718765&sr=8-1.

[26] J. Maynard Smith. Group Selection. *The Quarterly Review of Biology*, June 1976. ISSN 0033-5770. doi: 10.1086/409311. URL https://www.journals.uchicago.edu/doi/10.1086/409311. Publisher: Stony Brook Foundation, Inc.

[27] A. Gardner and A. Grafen. Capturing the superorganism: A formal theory of group adaptation. *Journal of Evolutionary Biology*, 22(4):659–671, 2009. ISSN 1010061X. doi: 10.1111/j.1420-9101.2008.01681.x.

[28] Boyd and Lorberbaum. No Pure Strategy Is Evolutionarily Stable in the Repeated Prisoners-Dilemma Game. *Nature*, 327(6117):58–59, 1987.

[29] Martin Nowak. Stochastic strategies in the Prisoner's Dilemma. *Theoretical Population Biology*, 38(1):93–112, August 1990. ISSN 00405809. doi: 10.1016/0040-5809(90)90005-G. URL https://linkinghub.elsevier.com/retrieve/pii/004058099090005G.

[30] Martin Nowak and Karl Sigmund. A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner's Dilemma game. *Nature*, 364(6432):56–58, July 1993. ISSN 0028-0836, 1476-4687. doi: 10.1038/364056a0. URL https://www.nature.com/articles/364056a0.

[31] MA Nowak and K Sigmund. Invasion Dynamics of the Finitely Repeated Prisoner's Dilemma. *Games and Economic Behavior*, 11:364–390, 1995. URL http://www.sciencedirect.com/science/article/pii/S089982568571055X.

[32] Julián García and Arne Traulsen. The structure of mutations and the evolution of cooperation. *PloS one*, 0(1):1–4, April 2012. ISSN 1932-6203. doi: 10.1371/journal.pone.0035287. URL http://dx.plos.org/10.1371/journal.pone.0035287.

[33] André. Mechanistic constraints and the unlikely evolution of reciprocal cooperation. *Journal of Evolutionary Biology*, 27(4):784–795, March 2014. ISSN 14209101. doi: 10.1111/jeb.12351. URL http://www.ncbi.nlm.nih.gov/pubmed/24618005. ISBN: 1010-061x.

[34] I Eshel. Evolutionary and continuous stability. *Journal of Theoretical Biology*, 103:99–111, 1983.

[35] J A J Metz. How Should We Define ‘ Fitness ’ for General Ecological Scenarios ? 1992.

[36] I Eshel, U Motro, and E Sansone. Continuous stability and evolutionary convergence. *Journal of theoretical biology*, 185(3):333–43, April 1997. ISSN 0022-5193. doi: 10.1006/jtbi.1996.0312. URL http://www.ncbi.nlm.nih.gov/pubmed/9156083.

[37] S. A.H. Geritz, É Kisdi, G. Meszéna, and J. A.J. Metz. Evolutionarily singular strategies and the adaptive growth and branching of the evolutionary tree. *Evolutionary Ecology*, 12(1):35–57, 1998. ISSN 02697653. doi: 10.1023/A:1006554906681.

[38] D. Waxman and S. Gavrilets. 20 Questions on Adaptive Dynamics. *Journal of Evolutionary Biology*, 18(5):1139–1154, September 2005. ISSN 1010-061X, 1420-9101. doi:

10.1111/j.1420-9101.2005.00948.x. URL https://academic.oup.com/jeb/article/18/5/1139-1154/7323785.

[39] John H. Gillespie. Molecular evolution over the mutational landscape. *Evolution*, 38(5): 1116–1129, September 1984. ISSN 0014-3820. doi: 10.1111/j.1558-5646.1984.tb00380.x. URL https://doi.org/10.1111/j.1558-5646.1984.tb00380.x.

[40] H Allen Orr. The Distribution of Fitness Effects Among Beneficial Mutations. 1526 (April):1519–1526, 2003.

[41] H. Allen Orr. The population genetics of beneficial mutations. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 365(1544):1195–1201, April 2010. ISSN 0962-8436, 1471-2970. doi: 10.1098/rstb.2009.0282. URL https://royalsocietypublishing.org/doi/10.1098/rstb.2009.0282.

[42] Rees Kassen and Thomas Bataillon. Distribution of fitness effects among beneficial mutations before selection in experimental populations of bacteria. *Nature Genetics*, 38(4):484–488, April 2006. ISSN 1061-4036, 1546-1718. doi: 10.1038/ng1751. URL https://www.nature.com/articles/ng1751.

[43] Adam Eyre-walker and Peter D Keightley. The distribution of fitness effects of new mutations. 8(August), 2007. doi: 10.1038/nrg2146.

[44] Alejandro Couce, Anurag Limdi, Melanie Magnan, Siân V. Owen, Cristina M. Herren, Richard E. Lenski, Olivier Tenaillon, and Michael Baym. Changing fitness effects of mutations through long-term bacterial evolution. *Science*, 383(6681):eadd1417, January 2024. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.add1417. URL https://www.science.org/doi/10.1126/science.add1417.

[45] J Maynard Smith and G A Parker. Logic of Asymmetric Contests. *Animal Behaviour*, 24(FEB):159–175, 1976.

[46] J. M. Smith. *Evolution and the theory of games*. Cambridge University Press, Cambridge, 1982. doi: 10.1016/0377-2217(83)90101-7. ISSN: 00030996.

[47] Ross Cressman. *Evolutionary Dynamics and Extensive Form Games*. MIT Press, 2003. ISBN 978-0-262-03305-3. Google-Books-ID: 187BOmBvyvQC.

[48] R Selten. Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games. *International Journal of Game Theory*, 4:25–55, 1975.

[49] R Selten. Evolutionary stability in extensive two-person games. *Mathematical Social Sciences*, 5:269–363, 1983.

[50] R Selten. Evolutionary stability in extensive two-person games - correction and further development. *Mathematical Social Sciences*, 16:223–266, 1988.

[51] John C Harsanyi. On the Rationality Postulates Underlying the Theory of Cooperative Games.

[52] R J Aumann and L S Shapley. Long-Term Competition: A game-theoretic analysis. In N Megiddo, editor, *Essays in Game Theory in Honor of Michael Maschler*. Springer, New York, 1994.

[53] Félix Geoffroy and André. The emergence of cooperation by evolutionary generalization. *Proceedings of the Royal Society B: Biological Sciences*, 288(1951):1–9, 2021. ISSN 14712954. doi: 10.1098/rspb.2021.0338.

[54] André. Contingency in the evolutionary emergence of reciprocal cooperation. *The American Naturalist*, 185(3):303–316, 2015. ISSN 00030147. doi: 10.1086/679625. URL http://www.journals.uchicago.edu/doi/10.1086/679625.

[55] Michael Taborsky, Joachim G. Frommen, and Christina Riehl. Correlated pay-offs are key to cooperation. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 371(1687):20150084, February 2016. ISSN 0962-8436, 1471-2970. doi: 10.1098/rstb.2015.0084. URL https://royalsocietypublishing.org/doi/10.1098/rstb.2015.0084.

[56] Redouan Bshary, Klaus Zuberbühler, and Carel P. Van Schaik. Why mutual helping in most natural systems is neither conflict-free nor based on maximal conflict. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 371(1687), 2016. ISSN 14712970. doi: 10.1098/rstb.2015.0091.

[57] Michael Taborsky, Michael A. Cant, and Jan Komdeur. *The Evolution of Social Behaviour*. Cambridge University Press, August 2021. ISBN 978-1-108-78863-2. Google-Books-ID: hMg8EAAAQBAJ.

[58] Richard C Connor. The benefits of mutualism: a conceptual framework. *Biological Reviews*, 70(3):427–457, 1995. Publisher: Wiley Online Library.

[59] Olof Leimar and Richard C RC Connor. By-product benefits, reciprocity, and pseudoreciprocity in mutualism. In Peter Hammerstein, editor, *Genetic and Cultural Evolution of Cooperation*, pages 203–222. The MIT Press, Cambridge MA, 2003. ISBN 0-262-08326-4. URL http://books.google.com/books?hl=en&lr=&id=aVh9jtWbG0wC&oi=fnd&pg=PA203&dq=By-product+benefits,+reciprocity,+and+pseudoreciprocity+in+mutualism&ots=guq9oHHgyj&sig=6wgXzqEU-g-zVp3qd8NwRSdNHD4.

[60] Richard Connor. Invested , extracted and byproduct benefits : A modified scheme for the evolution of cooperation. *Behavioural processes*, 2007. doi: 10.1016/j.beproc.2007.01.014. URL http://scholar.google.com/scholar?hl=

en&btnG=Search&q=intitle:Invested+,+extracted+and+byproduct+benefits+:+A+modified+scheme+for+the+evolution+of+cooperation#0.

[61] Olof Leimar and Peter Hammerstein. Cooperation for direct fitness benefits. *Philosophical Transactions of the Royal Society B-Biological Sciences*, 365(1553):2619–2626, 2010. doi: DOI10.1098/rstb.2010.0116. Publisher: ROYAL SOC Place: 6-9 CARLTON HOUSE TERRACE, LONDON SW1Y 5AG, ENGLAND.

[62] Richard Dawkins. *Climbing Mount Improbable*. Norton edition, 1996.

[63] Dilip Abreu and Ariel Rubinstein. The Structure of Nash Equilibrium in Repeated Games with Finite Automata. *Econometrica*, 56(6):1259, 1988. ISSN 00129682. doi: 10.2307/1913097.

[64] Banks and Sundaram. Repeated games, finite automata, and complexity. *Games and Economic Behavior*, 2(2):97–117, 1990. ISSN 10902473. doi: 10.1016/0899-8256(90)90024-O.

[65] David J. Cooper. Supergames Played by Finite Automata with Finite Costs of Complexity in an Evolutionary Setting. *Journal of Economic Theory*, 68(1):266–275, January 1996. ISSN 00220531. doi: 10.1006/jeth.1996.0015. URL https://linkinghub.elsevier.com/retrieve/pii/S0022053196900150.

[66] Oscar Volij. In Defense of DEFECT. *Games and Economic Behavior*, 39(2):309–321, May 2002. ISSN 08998256. doi: 10.1006/game.2001.0893. URL https://linkinghub.elsevier.com/retrieve/pii/S0899825601908930.

[67] Matthijs Van Veelen and Julian Garcıa. In and out of equilibrium II: evolution in repeated games with discounting and complexity costs. *Games and Economic Behavior*, 115:113–130, 2019.

[68] Martin A. Nowak and Karl Sigmund. Tit for tat in heterogeneous populations. *Nature*, 355(6357):250–253, 1992. ISSN 00280836. doi: 10.1038/355250a0.

[69] L Lehmann and L Keller. The evolution of cooperation and altruism – a general framework and a classification of models. *Journal of Evolutionary Biology*, 19:1365–1376, 2006.

[70] Brian Skyrms. *Evolution of the social contract*. Cambridge University Press, Cambridge, 1996. ISBN 978-1-139-92482-5. doi: 10.1017/CBO9781139924825. Publication Title: Evolution of the Social Contract.

[71] Richard C Connor. Altruism among non-relatives: alternatives to the 'Prisoner's Dilemma'. *Trends in Ecology & Evolution*, 10(2):84–86, 1995. ISSN 01695347. doi: 10.1016/S0169-5347(00)88988-0. Publisher: Elsevier.

[72] Olof Leimar and Peter Hammerstein. Evolution of cooperation through indirect reciprocity. *Proceedings of the Royal Society B: Biological Sciences*, 268(1468):745–753, 2001. ISSN 14712970. doi: 10.1098/rspb.2000.1573. Publisher: The Royal Society.

[73] N J Raihani and R Bshary. Resolving the iterated prisoner's dilemma: theory and reality. *Journal of Evolutionary Biology*, 24(8):1628–39, August 2011. ISSN 1420-9101. doi: 10.1111/j.1420-9101.2011.02307.x. URL `http://www.ncbi.nlm.nih.gov/pubmed/21599777`.

[74] André and Stefano Nolfi. Evolutionary robotics simulations help explain why reciprocity is rare in nature. *Scientific Reports*, in press:32785, 2016. ISSN 20452322. doi: 10.1038/srep32785. URL `http://www.nature.com/articles/srep32785`. Publisher: Nature Publishing Group.

[75] K. Ito, J. M. McNamara, A. Yamauchi, and A. D. Higginson. The evolution of cooperation by negotiation in a noisy world. *Journal of Evolutionary Biology*, 2017. ISSN 14209101. doi: 10.1111/jeb.13030.

# Supporting Information

## A Measuring mutational distance with finite automata

Here I seek a formal approach to obtaining a proxy for the mutational distance between two strategies. No approach is perfect, since the true mutational distance is a biochemical variable that cannot be measured from a simple model.

To this aim, I first characterize each strategy as a finite-state machine (see figures SI.1, SI.2, SI.3, SI.4, and SI.5). For the sake of conceptual clarity (and without losing generality), I always consider an initial state before the start of the game, called the initialization state, which is present in all games and all strategies.

To measure the distance between two strategies, I then have to solve an alignment problem. To do this, I try all possible alignments between the two strategies and define the distance between them as the smallest distance found in all these trials, i.e. with the best possible alignment. More precisely, I proceed as follows:

Let there be two strategies $S_1$ and $S_2$ with $k_1$ and $k_2$ states respectively (with $k_2 k_1$), and let there be an alignment $\mathcal{A}$ defined as a bijective mapping from the set of states of $S_1$ to the image of that set in the states of $S_2$ (i.e, each state of $S_1$ has one and only one image in $S_2$), subject only to the constraint that the image of the initialization state of $S_1$ must be the initialization state of $S_2$.

States in $S_2$ that are the image of a state in $S_1$ under alignment $\mathcal{A}$ are called homologous states according to $\mathcal{A}$, and transitions between two homologous states are called homologous transitions. States in $S_2$ that are not the image of any state in $S_1$ are called non-homologous states according to $\mathcal{A}$.

I then measure the following 4 quantities with this alignment:

- $\Delta_B$: number of homologous states where $S_2$ behaves differently from $S_1$.
- $\Delta_T$: number of homologous transitions with differing occurrence conditions between the two strategies.
- $\Delta_E$: number of non-homologous states in $S_2$ ($k_2 - k_1$).
- $\Delta_N$: number of transitions originating from non-homologous states in $S_2$.

The distance between $S_1$ and $S_2$ according to $\mathcal{A}$ is given by $d_{\mathcal{A}}(S_1, S_2) = \Delta_B + \Delta_T + \Delta_E + \Delta_N$. The distance between $S_1$ and $S_2$ is then obtained by minimizing $d_{\mathcal{A}}$ over all possible alignments, i.e., $d(S_1, S_2) = \min_{\mathcal{A}} d_{\mathcal{A}}(S_1, S_2)$. Figures SI.1, SI.2, SI.3, SI.4, and SI.5 show examples of this approach.

## A.1 Mutational distance is not the same as mathematical similarity

Many models in evolutionary game theory (e.g., Nowak and Sigmund (68), Lehmann and Keller (69)) implicitly adopt a different approach to measuring the distance between strategies.

They measure distances based on the implicit assumption that there is a direct correspondence between the formal mathematical description of strategies and the biological machinery that implements them, even though these are two entirely different things.

This can be illustrated with two examples.

In the ultimatum game, responder strategies are mathematically described as a demand level $q \in [0, 1]$. From a mathematical perspective, the two extreme strategies—one that accepts any offer ($q = 0$) and one that rejects all offers ($q = 1$)—are therefore closer to intermediate strategies, which accept some offers but reject others ($0 < q < 1$), than they are to each other.

In contrast, when distances are mesured using finite-state automata (see fig. SI.3), the two extreme strategies are represented by single-state automata, as they either accept or reject all offers unconditionally. Intermediate strategies, on the other hand, require two states, as they must condition their action on the offer they receive. As a result, the distance between the two extreme responder strategies is smaller than the distance between an extreme strategy and any intermediate strategy.

The same discrepancy between mathematical descriptions and finite automata is also observed in the repeated prisoner's dilemma with memory 1. In this game, strategies can be mathematically represented as a pair of actions $(X, Y)$, where $X$ denotes the action taken after the partner cooperates, and $Y$ denotes the action taken after the partner defects. Pure defection corresponds to $(D, D)$, pure cooperation to $(C, C)$, and tit-for-tat is represented as the intermediate pair $(C, D)$. Thus, from a mathematical perspective, Tit-for-tat appears to be an intermediate strategy between AllD and AllC.

In contrast, when distances are measured using finite-state automata, the opposite holds: tit-for-tat is farther from full defection than full cooperation is, as it introduces both an additional state with a new action and a new conditional transition between states (see fig. SI.5 for an illustration with the strategy Grim instead of tit-for-tat).

These two examples, the ultimatum game and the repeated prisoner's dilemma, illustrate the difference between the mathematical description of strategies and the characterization of a biological machine capable of implementing them. Modelers often make the implicit assumption that the mathematical representation of strategies should serve as a guide for determining their mutational distances. However, this is a mistake—one with significant consequences (32, 33).

From a biological mechanism perspective, accepting all offers or cooperating unconditionally simply requires a mechanism to acquire any available resource, or a constitutive mechanism that expresses cooperation at all times. Conversely, rejecting all offers or always defecting reflects a complete lack of interest in any resource, or in cooperation altogether.

In contrast, "intermediate" strategies are conditional strategies that require the ability to evaluate others' actions or offers and adjust responses accordingly. This demands an evolved mechanism specifically dedicated to executing such conditional behavior. The purpose of measuring distances based on finite-state automata is to reflect this notion.

# B Hawk-dove game

Consider a resource conflict game in which individuals compete for access to an indivisible resource, modeled using the classic hawk-dove framework. In this game, players invest in competition by choosing probabilities of escalating a conflict: they escalate (i.e., play Hawk) with a certain probability $p$ and refuse to escalate (i.e., play Dove) with the complementary probability $1 - p$. The payoffs for this interaction are outlined in Table SI.1. Players are also pre-assigned to distinct states by a move of Nature, and each strategy is therefore represented as a pair $(p_A, p_B)$, where $p_i \in [0, 1]$ denotes the probability of playing Hawk in state $i$.

| | Hawk | Dove |
|---|---|---|
| Hawk | P | T |
| Dove | S | R |

Table SI.1: **Payoff matrix of the hawk-dove game.** Parameters must respect $T > R > S > P$

## B.1 Mutational distances

As described in section A of this Supporting Information, finite-state automata are used to measure the distance between strategies (fig. SI.1). Two types of strategies are distinguished: symmetric strategies, which do not differentiate between states $A$ and $B$ and thus are, effectively, characterized by a single state automaton, and asymmetric strategies, which distinguish between $A$ and $B$, and therefore have two states, potentially with different probabilities, $p_A$ and $p_B$, of escalating in each.

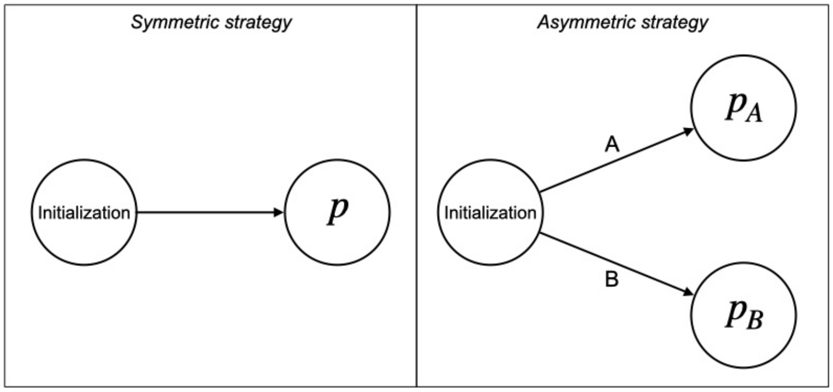

Figure SI.1: **Strategies in the hawk-dove game represented as finite automata**. The distance between two different symmetric strategies is 1. The distance between a symmetric and an asymmetric strategy is either 2 or 3, depending on whether the asymmetric strategy matches the symmetric in one of the two states or differs in both.

What matters for analyzing parsimonious dynamics are relative distances. Here, we see that the distance between a symmetric strategy, which expresses the same probability of escalating in both states, and an asymmetric strategy, which conditions its probability of escalating on the state, is always greater than the distance between two symmetric strategies. This reflects the idea that exhibiting asymmetric behavior in response to an arbitrary environmental feature does not occur spontaneously; it requires a biological mechanism capable of detecting and responding to the symmetry-breaking feature. Incorporating such a mechanism into a biological system demands a greater mutational change than any transition confined to symmetric strategies.

Consider an ancestral resident strategy $S$ that plays Hawk with a fixed probability $p \in [0, 1]$, independent of the state. Define a threshold value $\hat{p} \equiv \frac{T-R}{T+S-P-R}$. If the resident strategy is any $p \neq \hat{p}$, then any symmetric mutant strategy that plays Hawk with a probability $p'$ closer to $\hat{p}$ will be favored by selection. Asymmetric mutants—those employing different strategies depending on the state—may also be favored, but they are mutationally more distant from the resident strategy. As a result, under the constraint of parsimony, evolution will always preserve symmetry, leading to convergence toward the threshold value $\hat{p}$.

## B.2 Effects of ecological perturbations

Suppose now that a very small asymmetry exists between the two states. We illustrate this by assuming a slight asymmetry in the cost of conflict: players in state $A$ receive a payoff of $P + \epsilon$, while players in state $B$ receive $P - \epsilon$, instead of $P$, with $\epsilon \approx 0$.

Until the threshold value $\hat{p}$ is reached, this perturbation remains negligible, but this changes once $\hat{p}$ is attained. Against a symmetric resident strategy playing $\hat{p}$ in both states, all symmetric mutants—those adopting the same $p'$ in both states—are strictly neutral and therefore unable to invade. In contrast, asymmetric mutants that adjust their behavior based on the state can be favored. Specifically, a mutant that plays Hawk with a probability $\hat{p} + \delta$ in state $A$ and $\hat{p} - \delta$ in state $B$ gains a net advantage given by $\frac{T-R}{T+S-P-R}\epsilon\delta$. Thus, asymmetric mutants are favored as long as $\delta$ and $\epsilon$ have the same sign. If state $A$ provides some protection against the costs of conflict ($\epsilon > 0$), it becomes advantageous to play Hawk more frequently in this state. Crucially, these asymmetric mutants are the closest possible mutants able to invade, as the ability to condition behavior on the state is essential for securing an advantage in this scenario. Thus, the emergence of asymmetry here is consistent with parsimony.

Once symmetry breaking occurs and an asymmetric strategy becomes fixed, evolution amplifies the initial asymmetry. The resulting equilibrium is a fully state-dependent strategy, with players specializing their behavior based on their state. Starting from a state-independent ancestor $(p, p)$, evolution converges to one of two state-dependent equilibria: $(1, 0)$ or $(0, 1)$, depending on the sign of the small perturbation $\epsilon$. These outcomes correspond to the Bourgeois strategy in evolutionary biology MaynardSmith1976 and the concept of correlated equilibrium in game theory Aumann1974, AumannCorrelated1998.

The emergence of such conventions is parsimonious because even an infinitesimal quasi-arbitrary asymmetry can drive the evolution of genuinely state-dependent behavior. In other words, natural selection inherently favors conditioning actions on state, even in the absence of others doing so initially.

# C Divide-the-dollar game under simplified assumptions

Here, we aim to consider a situation where, unlike the hawk-dove game, asymmetry is not needed to resolve a coordination problem. Our goal is to illustrate as clearly as possible what happens in such a case.

To illustrate this, we consider a divide-the-dollar game. In this game, two individuals claim a fraction $q \in [0, 1]$ of a resource. If the combined claims do not exceed 1, both individuals receive exactly what they demand. However, if the total demand surpasses 1, neither individual receives anything, reflecting a sharp cost of competition.

For simplicity, we make two assumptions that streamline the analysis of evolutionary dynamics. First, we assume that the ancestral state always involves individuals making a demand $q_0 < 0.5$. Second, we assume that mutations affecting demand have only very small effects, in line with the standard assumptions of adaptive dynamics for a quantitative trait. We disregard the possibility of large-effect mutations. This latter assumption prevents equilibrium polymorphisms, which would complicate the analysis (70).

Like in the hawk-dove case above, at the start of the game, an arbitrary asymmetry is introduced through a move by Nature, assigning one individual to state $A$ and the other to state $B$. The index $i \in \{A, B\}$ is used to denote these two states, and a strategy is represented as a pair $(q_A, q_B)$, where $q_A$ and $q_B$ denote the demands made in states $A$ and $B$ respectively.

As long as the symmetry-breaking event is physically detectable by the players—even if it is entirely unrelated to their payoffs or competitive abilities—the game allows for an infinite range of asymmetric equilibria.

Any strategy profile where the player in state $A$ demands $q_A$ and the player in state $B$ demands $q_B = 1 - q_A$ constitutes a strict Nash equilibrium and an evolutionarily stable strategy. In a population where all individuals adopt this strategy, any mutant deviating from it is strictly selected against.

However, the principle of parsimony eliminates all these asymmetric equilibria, even in the presence of asymmetric perturbations, leaving only a single symmetric equilibrium.

## C.1 Mutational distances

Like in the hawk-dove game above (section B), finite-state automata are used to measure the distance between strategies (fig. SI.2). Two types of strategies are distinguished: symmetric strategies, which do not differentiate between states $A$ and $B$ and thus are, effectively, characterized by a single state automaton, and asymmetric strategies, which distinguish between $A$ and $B$, and therefore have two states, potentially with different demands, $q_A$ and $q_B$, in each. Like in the hawk-dove game, the distance between a symmetric strategy, which expresses the same demand in both states, and an asymmetric strategy, which conditions its demand on the state, is always greater than the distance between two symmetric strategies.

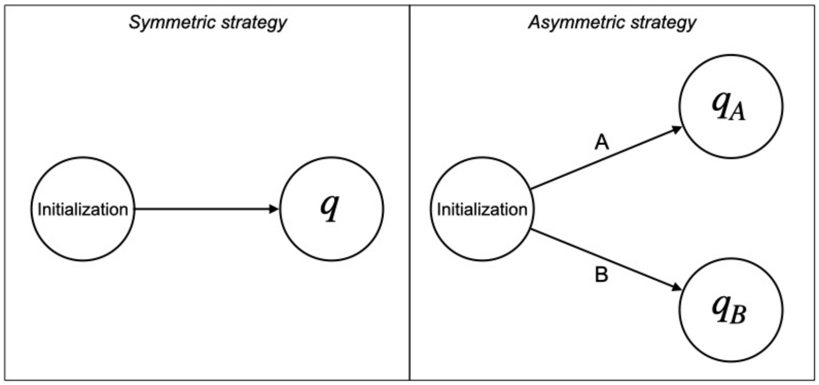

Figure SI.2: **Strategies in the divide-the-dollar game represented as finite automata**. The distance between two different symmetric strategies is 1. The distance between a symmetric and an asymmetric strategy is either 2 or 3, depending on whether the asymmetric strategy matches the symmetric in one of the two states or differs in both.

For any situation where the resident makes a demand $q < 0.5$, symmetric small-effect mutants that demand slightly more than the resident are always favored by selection. As a result, the evolutionary dynamics gradually push the population toward $q = 0.5$, which forms an equilibrium.

But starting from the same symmetric resident strategy with demand $q < 0.5$, selection also favors asymmetric mutants. For instance, selection would favor a mutant that demands a slightly higher $q'_A$ than the resident when in state $A$ while matching the resident's demand in state $B$. This asymmetry can then be further reinforced by selection.

For any asymmetric resident making two distinct demands, $q_A$ and $q_B$, as long as $q_A < 1$ and $q_B < 1 - q_A$, selection favors all small-effect mutants that demand strictly more than the resident in at least one state (while demanding at least as much in the other). That is, once symmetry is broken, evolutionary dynamics can act independently on the two states. Evolutionary dynamics can thus lead to any asymmetric equilibrium strategy characterized by $q_B = 1 - q_A$. This results in a large set of possible asymmetric equilibria, where both sides make complementary demands.

But the parsimony approximation eliminates all these asymmetric equilibria. Starting from a symmetric ancestral state with demand $q < 0.5$, an asymmetric mutant may be favored by selection. However, a symmetric mutant will always exist that is also favored by selection—and it is mutationally closer to the resident. As a result, the evolutionary transition from symmetry to asymmetry is not parsimonious. Thus, under the assumption of small mutational steps, the only parsimonious equilibrium from an ancestral state with demand $q < 0.5$ is the one where both players claim half of the resource.

### C.2 Effects of ecological perturbations

Suppose now that a slight perturbation generates an initial asymmetry between the two states. Here, a payoff asymmetry would obviously have no effect and is therefore not considered. Instead, we introduce an involuntary asymmetric perturbation in the demands of the two players. We consider an ancestral situation where both players have a genetically encoded target demand $q < 0.5$, but with slight unintended noise in their actual demands. Specifically, we assume that for any target demand $q$ in each state, a player in state $A$ always demands $q + \epsilon$, while a player in state $B$ always demands $q - \epsilon$, where $\epsilon$ is infinitesimally small ($\epsilon \approx 0$).

In this case, unlike the hawk-dove game above, deliberately adopting an asymmetric strategy offers no inherent advantage. Asymmetry does not solve a coordination problem that symmetry could not already resolve. Consequently, if evolution proceeds exclusively through parsimonious transitions, no symmetry breaking will occur in players' strategies. Any population starting with a symmetric demand will converge to the unique symmetric equilibrium, where the target demand is $q^* = 0.5$ in both states. In this equilibrium, individuals in state $A$ demand $0.5+\epsilon$, while those in state $B$ demand $0.5-\epsilon$. The equilibrium level of asymmetry thus precisely reflects the initial infinitesimal perturbation and is not amplified.

## D Non-credible threats, illustration with the ultimatum game

In the ultimatum game, an individual called the proposer unilaterally decides how to divide a resource, leaving the other participant, the responder, with no choice but to accept the offer or refuse it. If the responder rejects the offer, the entire interaction is canceled, resulting in no gain for either party. Since this form of punishment offers no benefit to the responder, it constitutes a non-credible threat.

In this game, a strategy profile is represented as a pair $(p, q)$, where $p$ denotes the offer made by the proposer, and $q$ represents the minimum offer required by the responder, below which the interaction is rejected. Hence, any strategy profile with $q > 0$ constitutes a non-credible threat, as it implies that the responder rejects certain offers without deriving any benefit from this refusal.

### D.1 An infinite range of equilibria involving non-credible threats

The only strategy that should represent a biologically reasonable equilibrium in this game is the profile $(0, 0)$, where the responder accepts any offer, no matter how small, and thus makes no non-credible threats.

Yet, all strategy profiles of the form $(p, p)$, where the proposer makes an offer $p > 0$ and the responder rejects any offer below $p$, are Nash equilibria and neutrally stable strategies (as there is no Evolutionarily stable strategy in this game). As long as the proposer consistently offers $p$, the responder's non-credible threat to reject offers below $p$ remains neutral, since it never needs to be enforced. Thus, this constitutes an equilibrium, as neither side has an incentive to change their strategy.

However, among all these equilibria, only the reasonable equilibrium $(0, 0)$, where the responder is willing to accept any offer, is a parsimonious equilibrium.

## D.2 Mutational distances

Finite-state automata are used to measure the distance between responder strategies (fig. SI.3). Two types of responder strategies are distinguished: (i) Unconditional responders, who do not adjust their response based on the offer they receive. They either accept or reject all offers regardless of their value and are thus represented by a single-state automaton. (ii) Conditional responders, who modify their response depending on the offer. They can either accept or reject based on its value, requiring a two-state automaton.

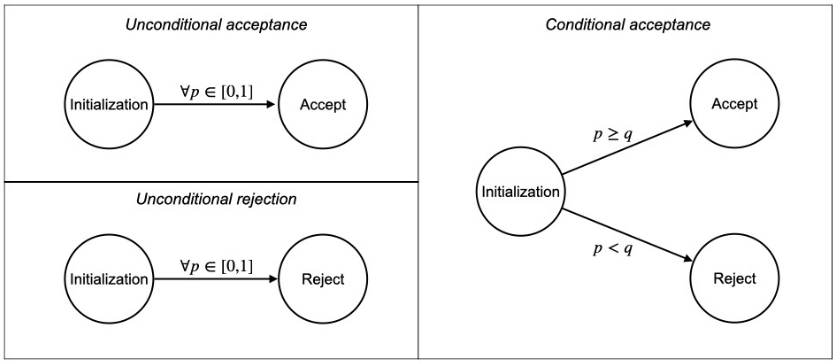

Figure SI.3: **Strategies in the ultimatum game represented as finite automata**. The distance between unconditional acceptance and unconditional rejection is 1. The distance between either unconditional acceptance or unconditional rejection and conditional acceptance is 2.

What matters for characterizing parsimonious dynamics are relative distances. Here, the distance between the two extreme responder strategies—those that either accept or reject all offers—is smaller than the distance between an extreme strategy and any intermediate strategy, which accepts some offers while rejecting others.

## D.3 A single parsimonious equilibrium free of non-credible threats

Consider an ancestral strategy profile where the proposer offers any $p_0 > 0$, and the responder is unconditional. If the ancestral responder unconditionally accepts all offers, a mutant who accepts offers only above a threshold $x > 0$ would be at best neutral (if $x < p_0$) or counter-selected (if $x \geq p_0$). Thus, a transition to conditional acceptance is not feasible in this case. If the ancestral responder unconditionally rejects all offers, a conditional mutant responder who accepts only offers above a threshold $x$ could potentially be favored if $x \leq p_0$. However, a non-conditional mutant responder who accepts all offers regardless of their value would be equally favored and is assumed to be mutationally closer to the resident strategy, as it requires no conditional mechanism. Therefore, a transition to conditional acceptance can never occur parsimoniously.

Starting from any unconditional ancestral strategy profile where the proposer offers $p_0 > 0$, the only parsimonious equilibrium is one where the proposer offers the minimum possible amount ($p_0 = 0$) and the responder accepts any offer.

Note that an unconditional strategy profile where the proposer offers $p_0 = 0$, and the responder either accepts or rejects unconditionnally all offers is also a Nash equilibrium. However, this equilibrium cannot be reached from any other starting point.

# E Investment game

The investment game is an asymmetric repeated game played by two individuals, called the Investor and the Responder. The repeated game is preceded by a move by Nature that randomly determines the type of stage game that the two individuals will then play repeatedly.

Three types of stage games are possible: (1) The stage game can be a by-product cooperation game, in which both the responder and the investor automatically benefit from the investment. (2) The stage game can be a trust game, in which the investment has a net cost to the investor, but the responder can reward him afterwards and still make a net profit. (3) The stage game can be an inefficient cooperation game, in which the investment has a net cost to the investor and the responder cannot reward him in a mutually beneficial way. These three games differ in their payoffs, but they are identical in terms of their action sets.

Regardless of the game type chosen by Nature, the investor has two possible actions in each round: *Invest* or *Decline.* Then, if the investor chooses to invest in a given round, the responder has two possible actions: *Reward* or *Defect.* If the investor chooses to decline, then the responder has no decision to make.

In each round of the game, four outcomes are possible: (1) "Invest-No-Reward," where the investor invests but receives no reward, neither from their partner nor from Nature; (2) "Invest-Natural-Reward," where the investor invests and receives a reward from Nature (applicable only in byproduct games); (3) "Invest-Social-Reward," where the investor invests and is rewarded by their partner; and (4) "Decline," where the investor chooses not to invest in that round.

The payoffs of each stage game are as follows:

**Trust game**: The cost of investing is $c$, the benefit of receiving the investment is $b$, the cost of rewarding is $r$, and the benefit of receiving the reward is also $r$ (i.e., the reward is a conservative transfer), and we assume that we have the relation $b > r > c > 0$.

**By-product game**: The natural reward received after investing is $a > c$, (hence the net benefit is $a - c > 0$) and all other parameters are as in the trust game.

**Inefficient cooperation game**: For simplicity and without loss of generality, I assume that the benefit of receiving the investment is 0, and that all other parameters are as in the trust game.

The same stage game is assumed to repeat indefinitely, with a constant stopping probability $\delta$ between any two rounds. The expected payoff in round $n \geq 0$ is therefore discounted by a factor of $(1-\delta)^n$. By convention, the first round is designated as $n = 0$. If an individual earns $x$ per round, their total payoff over the entire interaction is given by $x/\delta$.

Five additional assumptions are made:

- Only the responder knows the type of game being played, capturing the essence of why reputation matters: the investor must infer their partner's payoffs by observing their behavior.
- The investor can distinguish between the two types of rewards they might receive: natural rewards, automatically obtained in the by-product cooperation game, and social rewards, provided by trustworthy partners in the trust game.
- Both players are aware of the round number, but only starting from the first investment made by the investor; this initial investment triggers the round number counter.
- Parameters are assumed such that $\frac{b-r}{\delta} > b$. In other words, in the trust game, it is always worthwhile for the responder to reward their partner to continue the interaction and earn $b - r$ in future rounds.
- Parameters are also assumed to satisfy $\frac{a-c}{\delta} p_b > -c(p_t + p_i)$. This means that, even in the worst-case scenario, it is always beneficial to test the game in the first round to determine whether it is a by-product game, which would allow securing $a - c$ in future rounds.

## E.1 Definition of strategy profiles

We begin by defining three types of strategy profiles that will be central to the reasoning that follows. See figure SI.4 for their representation as finite automata.

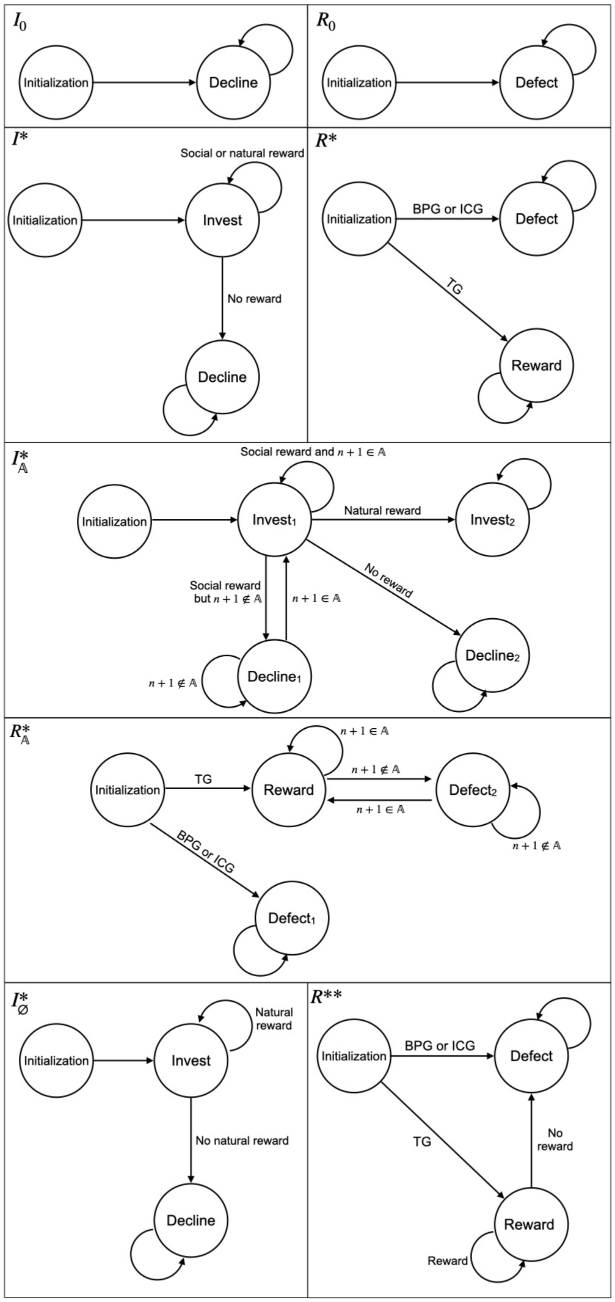

Figure SI.4: **Strategies in the investment game represented as finite automata**. The distances are provided in Tables SI.4 and SI.5 of this supporting information. TG: Trust game; BPG: By-product cooperation game; ICG: Inefficient cooperation game.

### E.1.1 Unconditional defection

The profile $S_0 = (I_0, R_0)$ represents pure, unconditional defection. The investor's strategy, $I_0$, consists of playing *Decline* in all situations, while the responder's strategy, $R_0$, consists of playing *Defect* in all situations.

### E.1.2 Unconditional cooperation

The profile $\tilde{S} = (\tilde{I}, \tilde{R})$ represents pure, unconditional cooperation. The investor's strategy, $\tilde{I}$, consists of playing *Invest* in all situations, while the responder's strategy, $\tilde{R}$, consists of playing *Reward* in all situations.

### E.1.3 Conditional cooperation

Consider any subset of strictly positive integers $A \subseteq Z^+$, and consider the strategy profile $S^*_A = (I^*_A, R^*_A)$ defined as follows.

Note: for simplicity, this definition does not account for the strategy's response to its own deviations, as we are not concerned with subgame perfection (see section E.6 for a discussion on this topic). The full description of the strategy is more precisely captured in its finite automaton (figure SI.4).

**Investor's strategy $I^*_A$:** - Play *Invest* in any round $n$ iff at least one of the following conditions is met:

- $n = 0$ (i.e., this is the first round of the game).
- The investment in round 0 resulted in a natural reward.
- $n \in A$ and all past investments in rounds $m \in A \cup \emptyset$ resulted in a social reward.

- Play *Decline* otherwise.

**Responder's strategy $R^*_A$:**

- Play *Reward* iff current round number $n \in A \cup \emptyset$, and the game is a trust game.
- Play *Defect* otherwise.

Simply put, the investor's strategy, $I^*_A$, consists of testing the game once to determine the nature of the rewards received (natural, social, or absent). The investor then continues to invest in all rounds if a natural reward is received, invests only in rounds belonging to $A$ if the reward is social, and stops investing entirely if no reward is received. The responder's strategy, $R^*_A$, involves actively rewarding investments only if the game is a trust game and the round number belongs to $A \cup \emptyset$ .

We will also consider the two extreme profiles within this family, corresponding to $A = \emptyset$ and $A = Z^+$, as they represent important strategies.

The profile $S^*_{Z+}$, which we denote more simply as $S^*$, represents plain conditional cooperation. In simple terms, the investor's strategy consists of testing the game once to determine whether a reward is received. If any reward is given, whether natural or social, the investor continues to invest in all subsequent rounds. However, if no reward is received, investment stops permanently. The responder's strategy consists of actively rewarding all investments if the game is a trust game.

The profile $S^*_\emptyset$ represents byproduct cooperation. Here, the investor tests the game once to determine the nature of the rewards. If a natural reward is received, investment continues in all rounds; otherwise, investment ceases entirely. The responder never rewards and always plays *Defect*.

## E.2 Payoffs

We define $P(X, Y)$ as the payoff of an individual playing strategy $X$ against a partner playing strategy $Y$. For example, $P(I_1, R_2)$ represents the payoff of an investor using strategy $I_1$ against a responder using strategy $R_2$, while $P(R_1, I_2)$ represents the payoff of a responder using strategy $R_1$ against an investor using strategy $I_2$.

To calculate the payoffs of individuals following a profile $S^*_A$, we define $\mathbf{1}_A(n)$ as the indicator function of $A$, such that:

$$\mathbf{1}_A(n) \equiv \begin{cases} 1 & \text{if } n \in A, \\ 0 & \text{if } n \notin A. \end{cases} \tag{2}$$

The effective density of $A$ after round $n$ is defined as:

$$\rho_A(n) \equiv \sum_{k=n+1}^{+\infty} \mathbf{1}_A(k)(1-\delta)^{k-n} \tag{3}$$

The payoffs of individuals following the strategy profile $S^*_A$ are therefore as follows:

$$P(I^*_A, R^*_A) = p_t(r-c)\left(1+\rho_A(0)\right) + p_b\frac{(a-c)}{\delta} - p_i c \tag{4}$$

$$P(R^*_A, I^*_A) = p_t(b-r)\left(1+\rho_A(0)\right) + p_b\frac{b}{\delta} \tag{5}$$

Beyond the profile $S^*_A$ , all other usefull payoffs are shown in Tables SI.2 and SI.3.

| | $R_0 = R^*_\emptyset$ | $R^*$ |
|---|---|---|
| $I_0$ | $0$ | $0$ |
| $\tilde{I}$ | $\frac{ap_b - c}{\delta}$ | $p_b\frac{a-c}{\delta} + p_t\frac{r-c}{\delta} - p_i\frac{c}{\delta}$ |
| $I^*_\emptyset$ | $p_b\frac{a-c}{\delta} - (p_i + p_t)c$ | $p_b\frac{a-c}{\delta} + p_t(r - c) - p_i c$ |
| $I^*$ | $p_b\frac{a-c}{\delta} - (p_i + p_t)c$ | $p_b\frac{a-c}{\delta} + p_t\frac{r-c}{\delta} - p_i c$ |

Table SI.2: **Payoff matrix for the Investor in the investment game**. Each cell shows the Investor's payoff in interactions between a given pair of strategies.

| | $I_0$ | $I^*$ |
|---|---|---|
| $R_0 = R^*_\emptyset$ | $0$ | $p_b\frac{b}{\delta} + p_t b$ |
| $\tilde{R}$ | $0$ | $(p_b + p_t)\frac{b-r}{\delta} - p_i\frac{r}{\delta}$ |
| $R^*$ | $0$ | $p_b\frac{b}{\delta} + p_t\frac{b-r}{\delta}$ |

Table SI.3: **Payoff matrix for the Responder in the investment game**. Each cell shows the Responder's payoff in interactions between a given pair of strategies.

## E.3 An infinite range of equilibria

Consider any subset of natural numbers $A \subset Z^+$, such that:

$$\forall n0, \rho_A(n) > \frac{r}{b-r} \tag{6}$$

In words, we consider all subsets $A$ whose density is sufficiently high within $Z^+$ to ensure that it is always beneficial for a responder to continue the interaction, regardless of the round they are in. This includes a wide range of subsets of $Z^+$ with varying densities, from $Z^+$ itself, with a density $\rho_{Z^+}(n) = \frac{1-\delta}{\delta}$, down to the lowest possible density, $\frac{r}{b-r}$.

Here, I aim to show that, provided condition 6 is satisfied, any strategy profile $S^*_A = (I^*_A, R^*_A)$ is a Nash equilibrium. To do so, I examine all possible deviations on both sides that actually occur along the equilibrium path of the strategy and demonstrate that each of them is strictly disadvantageous. Other deviations, which never occur against a partner playing $S^*_A$, are neutral.

### E.3.1 Deviations on the investor's side

Against a responder playing strategy $R^*_A$, the investor can find themselves in only five possible situations. We will examine the effects of deviating in each case.

1. First round ($n = 0$): Nothing has happened yet, so the investor must test the game immediately; otherwise, they waste time.

2. Any round $n > 1$, where the event "Invest-No-reward" occurred in the first round: The investor is in an inefficient cooperation game. They should decline; otherwise, they incur an unnecessary cost.

3. Any round $n > 1$, where the event "Invest-Natural reward" occurred in the first round: The investor is in a byproduct cooperation game. They should invest; otherwise, they miss a profitable opportunity.

4. A round $n > 1$ with $n \in A$, where the event "Invest-Social reward" occurred in the first round: The investor is in a trust game. They should invest; otherwise, they miss a profitable opportunity.

5. A round $n > 1$ with $n \notin A$, where the event "Invest-Social reward" occurred in the first round: The investor should decline; otherwise, they incur an unnecessary cost.

All other possible states are not reached at equilibrium against $R^*_A$. In particular, a situation where the game is a trust game but the responder has defected once in a round belonging to $A$ does not occur along the equilibrium path.

We now examine each of the cases above and formally show that no deviation is beneficial:

1. In the first round, the strategy $I^*_A$ prescribes investing to test the game, and the continuation payoff for the investor is then given by $P(I^*_A, R^*_A)$ in equation 4 above. If the investor deviates by declining in the first round and then returns to the resident strategy, the first game test would only begin in the next round, delaying the increment of the round number as well. Consequently, the continuation payoff remains the same as that of the resident but shifted one round later, meaning it is discounted by a factor of $1 - \delta < 1$, making the deviation strictly costly.

2. If the investor received no reward in the first round, then the game is an innefficient cooperation game. In this case, for all subsequent rounds $n$, $I^*_A$ prescribes declining, and the investor's continuation payoff after $n$ is zero. If the investor deviates by investing instead, their continuation payoff becomes strictly negative ($-c$), making the deviation strictly costly.

3. If the investor received a natural reward in the first round, then the game is a byproduct cooperation game. In this case, for all subsequent rounds $n$, $I^*_A$ prescribes investing, and the investor's continuation payoff after $n$ is $\frac{a-c}{\delta}$. If the investor deviates by declining at round $n$ and then resumes the $I^*_A$ strategy, they forgo one round of gain, resulting in a continuation payoff of $\frac{a-c}{\delta}(1-\delta)$, which is strictly lower.

4. If the investor received a social reward in the first round, and the current round is $n \in A$, the game is a trust game, and the responder will reward this round. If the investor follows $I^*_A$ and invests, their continuation payoff is $(r-c)(1+\rho_A(n))$, as they gain $r-c$ now plus future payoffs. If they deviate by declining, they miss an opportunity to make a gain,

resulting in a continuation payoff of only $(r - c)\rho_A(n)$. The deviation is therefore strictly costly.

5. If the investor received a social reward in the first round but the current round is $n \notin A$, the game is a trust game, but the responder will not reward this round. If the investor follows $I^*_A$ , they must decline, and their continuation payoff is $(r-c)\rho_A(n)$. If they deviate by investing, they pay an unnecessary immediate cost, $c$, while obtaining the same continuation payoff afterward. This deviation is thus strictly costly.

### E.3.2 Deviations on the responder's side

Against an investor playing the strategy $I^*_A$, the responder faces only three possible situations in which they must make a decision.

1. If the game is a by-product cooperation game and the round number is any $n \geq 0$, $R^*_A$ prescribes playing *Defect*. Deviating cannot be beneficial because rewarding has no effect on the investor's behavior in this case and only incurs a cost.

2. If the game is an inefficient cooperation game and the round number is $n = 0$, $R^*_A$ prescribes playing *Defect*. In this case, the responder's continuation payoff is 0, as they gain no benefit from investments. If the responder deviates once by rewarding the first investment and then returns to $R^*_A$, they will have incurred a net cost, $r$.

3. If the game is a trust game and the round number is $n \in A$, $R^*_A$ prescribes playing Reward after this investment. The responder's continuation payoff is then $(b-r)(1+\rho_A(n))$, as they receive an immediate benefit of $b-r$ plus their future gains. If instead, the responder deviates and plays Defect in this round, they gain an immediate profit by keeping $r$, but this ends the game, leaving them with a continuation payoff of only $b$. Under assumption 6, $b$ is always strictly lower than the resident's continuation payoff, meaning the future value of the game is always sufficient to make rewarding worthwhile in every round.

In conclusion, all deviations along the equilibrium path are strictly costly, while other deviations occurring off the equilibrium path are neutral. For any set of natural numbers $A$ that satisfies condition 6, the strategy profile $S^*_A$ is a Nash equilibrium.

Notably, among many others, this includes the most cooperative profile, $S^*_{Z+} = S^*$, which is also a Nash equilibrium.

### E.3.3 The byproduct cooperation profile, $S^*_\emptyset$ , is also a Nash equilibrium

The proof is slightly different in the particular case where $A = \emptyset$.

On the investor's side:

In the first round, if the investor plays *Invest* as prescribed by $I^*_\emptyset$, their expected payoff over the entire game is $-c(p_t + p_i) + \frac{a-c}{\delta} p_b$, which, by assumption, is strictly greater than 0. This means the parameters are such that it is always worth testing the game once to determine whether it is a by-product game. On the other hand, if the investor deviates in this round and instead tests the game for the first time in the next round, they receive an immediate payoff of 0 and then obtain the same continuation payoff, but discounted by a factor of $1 - \delta$. Thus, postponing the first investment is strictly costly at all rounds.

In subsequent rounds, if the investor received a natural reward in the first round, $I^*_\emptyset$ prescribes investing, which yields the continuation payoff $(a-c)/\delta$. If the investor deviates, they will only postpone this payoff and receive $(a-c)(1-\delta)/\delta$ instead, which is strictly lower.

On the other hand, if the investor received no reward in the first round, $I^*_\emptyset$ prescribes declining, which results in a continuation payoff of 0. If the investor instead chooses to invest, they incur an immediate cost of $-c$ with no future benefit. Hence, this deviation is also strictly costly.

On the responder's side, $R^*_\emptyset$ stipulates to defect in all cases and in all rounds. If the investor plays $I^*_\emptyset$ then regardless of the nature of the game, his behavior is independent of the responder's actions, so there is never any benefit for the responder to reward him. Hence any deviation from $R^*_\emptyset$ is strictly costly.

In conclusion, all deviations on the equilibrium path are strictly costly. Other deviations, occuring out of equilibrium are neutral. The strategy $S^*_\emptyset$ is a Nash equilibrium.

In conclusion, there is a wide variety of strategies that constitute Nash equilibria in this game, as in all repeated games. But we will now show that the parsimony criterion eliminates the vast majority of these equilibria. To do so, we will first demonstrate that the two extreme equilibria—the fully cooperative equilibrium $S^*$ and the by-product cooperation equilibrium $S^*_\emptyset$—are parsimonious equilibria from $S_0$. Then, we will show that none of the intermediate Nash equilibria in the family $S^*_A$, with $A \neq \emptyset$ and $AZ^+$, are parsimonious from $S_0$.

A point of terminology: when an evolutionary transition from a strategy profile $S$ to a strategy profile $S'$ is parsimonious (as defined in section 2), we say that $S'$ is parsimoniously evolvable in one step from $S$. And, more generally, if there exists at least one trait substitution sequence from $S$ to $S'$ that consists only of parsimonious transitions, we say that $S'$ is parsimoniously evolvable (possibly in multiple steps) from $S$.

## E.4 The two extreme equilibria are evolutionarily parsimonious from the ancestral strategy $S_0$

Gist of the proof: $S^*_\emptyset$ and $S^*$ are the closest possible conditional strategy profiles to the ancestral profile $S_0$ from a mutational perspective. Since there is a direct selection pressure favoring conditional investments, these two strategies can invade a non-cooperative resident population and are thus parsimoniously evolvable from $S_0$.

Here is the detailed proof:

### E.4.1 Mutational distances

Based on cellular automata, as explained in section A, the mutational distances between strategies can be calculated and are presented in Tables SI.4 and SI.5.

| | $I_0$ | $\tilde{I}$ |
|---|---|---|
| $I_0$ | 0 | 1 |
| $\tilde{I}$ | 1 | 0 |
| $I^*_\emptyset$ | 4 | 3 |
| $I^*$ | 4 | 3 |
| $I^*_A$ | 9 | 8 |

Table SI.4: **Mutational distances in the investment game on the investor's side.**

| | $R_0 = R^*_\emptyset$ | $\tilde{R}$ |
|---|---|---|
| $R_0 = R^*_\emptyset$ | 0 | 1 |
| $\tilde{R}$ | 1 | 0 |
| $R^*$ | 3 | 2 |
| $R^*_A$ | 8 | 7 |

Table SI.5: **Mutational distances in the investment game on the responder's side.**

### E.4.2 The fully cooperative strategy $S^*$ is an evolutionarily parsimonious equilibrium from the ancestral strategy $S_0$

To show that the strategy profile $S^* = (I^*, R^*)$ is parsimoniously evolvable from the ancestral strategy $S_0 = (I_0, R_0)$, I will first show that the strategy $(I^*, R_0)$ is parsimoniously evolvable from $(I_0, R_0)$, and then show that the strategy $(I^*, R^*)$ is parsimoniously evolvable from $(I^*, R_0)$.

Let us start with showing that $(I^*, R_0)$ is parsimoniously evolvable from $S_0$.

First, under our assumptions, we always have $P(I^*, R_0) > P(I_0, R_0)$ (see payoff table SI.2), meaning that the conditional investment strategy $I^*$ is always strictly better than the ancestral strategy $I_0$ in an environment composed of $R_0$. This is because the long-term benefit of by-product games always outweighs the risk of testing the games once.

However, for $I^*$ to be parsimoniously evolvable from $I_0$, it is not enough for it to be capable of invasion. There must also be no other strategy strictly closer to $I_0$ that can invade as well. For instance, this could be the case for the unconditional investment strategy $\tilde{I}$, which is strictly closer to $I_0$ (see distance table SI.4).

Here, we must distinguish between two cases.

If $p_b a < c$, then $P(\tilde{I}, R_0) < 0$ (see payoff table SI.2), meaning that the unconditional investment strategy $\tilde{I}$ performs worse than $I_0$ in an environment composed of responders $R_0$, who never reward. This is because the cost incurred in non-cooperative games is too high, making blind investment unprofitable on average. In this case, $I^*$ is the closest strategy to $I_0$ that can successfully invade (see distance table SI.4). Therefore, the strategy profile $(I^*, R_0)$ is parsimoniously evolvable from $(I_0, R_0)$ in one step.

If $p_b a > c$, on the other hand, then $P(\tilde{I}, R_0) > 0$ (see payoff table SI.2), meaning that the unconditional investment strategy $\tilde{I}$ performs better than $I_0$ in an environment composed of $R_0$. This is because the benefit from by-product games is sufficient to compensate for the cost. In this case, $I^*$ is not parsimoniously evolvable from $I_0$ in one step. Instead, $\tilde{I}$ is the strategy that is parsimoniously evolvable from $I_0$ in one step.

However, in this case, $I^*$ is parsimoniously evolvable in two steps. Once $\tilde{I}$ has fixed in the population, $I^*$ becomes parsimoniously evolvable. Indeed, we have $P(I^*, R_0) > P(\tilde{I}, R_0)$ (see payoff table SI.2), since $I^*$ introduces an adaptive refinement in the investment strategy. Moreover, $I^*$ is the closest possible strategy to $\tilde{I}$ that can successfully invade (see distance table SI.4).

Thus, in all cases, the strategy profile $(I^*, R_0)$ is parsimoniously evolvable from the ancestral strategy $S_0$, either in one step or in two steps.

Let me now show that the fully conditional profile $S^* = (I^*, R^*)$ is parsimoniously evolvable from $(I^*, R_0)$.

First of all, under our assumptions, the parameters are such that we always have $P(R^*, I^*) > P(R_0, I^*)$ (see payoff table SI.3). This means that $R^*$ can always replace $R_0$ in an environment composed of $I^*$, since it benefits from trust games, whereas $R_0$ does not.

However, we face the same issue as before: for $R^*$ to be parsimoniously evolvable from $R_0$, it is not enough that it can invade—it must also be the closest strategy to $R_0$ that can do so. Otherwise, a strategy even closer to $R_0$ might invade first. In this case, the unconditional reward strategy $\tilde{R}$ could be such a candidate, as it is strictly closer to $R_0$ (see distance table SI.5).

Thus, once again, we must distinguish between two cases.

If $r > bp_t(1 - \delta)$, then $P(\tilde{R}, I^*) < P(R_0, I^*)$ (see payoff table SI.3), meaning that the unconditional reward strategy $\tilde{R}$ is not better than $R_0$ in an environment composed of $I^*$. This is because the cost of rewarding all games indiscriminately is too high compared to the benefit. In this case, $R^*$ is the closest strategy to $R_0$ that can successfully invade (see distance table SI.5). Therefore, the strategy profile $(I^*, R^*)$ is parsimoniously evolvable from $(I^*, R_0)$ in one step.

If $r < bp_t(1-\delta)$, on the other hand, then $P(\tilde{R}, I^*) > P(R_0, I^*)$ (see payoff table SI.3), meaning that the unconditional reward strategy $\tilde{R}$ performs better than $R_0$ in an environment composed of $I^*$. This is because rewarding all games indiscriminately is worthwhile in this case. As a result, it is not $R^*$ but rather $\tilde{R}$ that is parsimoniously evolvable from $R_0$ in one step.

However, in this case, $R^*$ is still parsimoniously evolvable in two steps. Once $\tilde{R}$ has fixed in the population, $R^*$ becomes parsimoniously evolvable. We always have $P(R^*, I^*) > P(\tilde{R}, I^*)$ (see payoff table SI.3), since $R^*$ introduces an adaptive refinement in the reward strategy. Moreover, $R^*$ is the closest possible strategy to $\tilde{R}$ that can successfully invade (see distance table SI.5).

Thus, in all cases, the strategy profile $(I^*, R^*)$ is parsimoniously evolvable from the strategy profile $(I^*, R_0)$, either in one step or in two steps.

We have therefore shown that, under our assumptions, the strategy profile $(I^*, R_0)$ is always parsimoniously evolvable from the ancestral strategy $S_0$, and that the profile $S^*$ is, in turn, parsimoniously evolvable from $(I^*, R_0)$. As a consequence, the conditional cooperation profile $S^*$ is parsimoniously evolvable from the ancestral profile $S_0$. Since $S^*$ is also a Nash equilibrium, it is an evolutionarily parsimonious equilibrium from $S_0$.

### E.4.3 The byproduction cooperative equilibrium $S^*_\emptyset$ is also evolutionarily parsimonious from $S_0$

$S^*_\emptyset$ is the least cooperative equilibrium, supporting cooperation only in by-product games and never in trust games.

First, note that the responder's strategy in $S^*_\emptyset$ is exactly the same as in the ancestral strategy $S_0$, meaning $R^*_\emptyset = R_0$. This strategy simply consists of never actively rewarding any partner. Hence, to show that $S^*_\emptyset$ is parsimoniously evolvable from $S_0$, we only need to show that the strategy $I^*_\emptyset$ is parsimoniously evolvable from $I_0$ in an environment where the responder plays $R_0$.

In an environment where the responder plays $R_0$, the strategy $I^*_\emptyset$ behaves exactly like $I^*$, and its payoff is therefore identical, i.e., $P(I^*_\emptyset, R_0) = P(I^*, R_0)$ (see payoff table SI.2). Furthermore, the distances between $I^*_\emptyset$ and the two other strategies ($I_0$ and $\tilde{I}$) are the same as the distances between $I^*$ and these two strategies. As a result, the proof given above showing that $I^*$ is parsimoniously evolvable from $I_0$ applies in exactly the same way to $I^*_\emptyset$.

Hence, like $I^*$, the strategy $I^*_\emptyset$ is parsimoniously evolvable from $I_0$ in an environment where the responder plays $R_0$.

As a result, the strategy profile $S^*_\emptyset$ is parsimoniously evolvable from the ancestral profile $S_0$. Since $S^*_\emptyset$ is also a Nash equilibrium, it is a parsimonious equilibrium from $S_0$.

## E.5 Intermediate equilibria of the family $S_A^*$ are not parsimoniously evolvable from $S_0$

We refer to intermediate equilibria as all equilibrium strategies that invest in some rounds of the trust game but not in all rounds, i.e., strategies belonging to the family $S_A^*$, where $A$ is a non-empty subset of $Z^+$. We will also refer to these as round number-dependent strategies.

Additionally, for the investor, we use the term reward-dependent to describe strategies in which the investor tests the games and continues investing only when they yield a given reward, either natural or social, or both. Investors in intermediate equilibria are thus both round-number dependent and reward-dependent. In contrast, the ancestral non-conditional strategy $I_0$ is neither round-number dependent nor reward-dependent.

For the responder, we use the term game-dependent to describe strategies in which the responder rewards some games (typically the trust game) but not all. Responders in intermediate equilibria are thus both round-number dependent and game-dependent. In contrast, the ancestral non-conditional strategy $R_0$ is neither round-number dependent nor game-dependent.

Here, we will show that a round-number dependent strategy profile can never be evolutionarily parsimonious from a non-conditional ancestral strategy profile.

Here is first an intuitive verbal explanation:

To evolutionarily reach a round-number dependent strategy from an ancestral round-number independent strategy, there must be a point where round-number dependence appears on one side (investor or responder), despite being initially absent on the other side. In either case, in a situation where the other side is not round-number dependent, the very ability to be dependent on the round number cannot provide a strictly positive advantage.

Mutationally speaking, a round-number independent strategy is always closer to an ancestral strategy that is itself round-number independent than a round-number dependent strategy is (see distance table).

Consequently, there is no scenario in which a transition from a round-number independent to a round-number dependent strategy is parsimonious.

To be more precise, we will now examine each of the possible scenarios in detail.

To evolutionarily reach a round-number dependent strategy from a round-number independent ancestral strategy, there are only two possible scenarios. (1) A round-number dependent strategy on the investor's side invades an environment where the responder is not round-number dependent. (2) A round-number dependent strategy on the responder's side invades an environment where the investor is not round-number dependent.

### E.5.1 Round-number dependence on the investor's side first

Starting from the ancestral strategy $S_0 = (I_0, R_0)$, the only scenario in which round-number dependence first appears on the investor's side is one where this conditionality successfully invades in an environment composed of $R_0$ (i.e., responders who do not reward any game).

To analyze this, we must consider two possible cases:

(i) If $ap_b < c$, then an investor who is not reward-dependent cannot invade. In this case, the only possible way to reach a round-number dependent strategy is through the invasion of a strategy $I^*_A$ that is both round-number dependent and reward-dependent.

However, we know that in this situation, an investor who is reward-dependent but not round-number dependent (i.e., $I^*$ or $I^*_\emptyset$) can also invade while being strictly closer to the resident strategy. This means that $I^*_A$ is not parsimoniously evolvable in one step. Furthermore, once $I^*$ or $I^*_\emptyset$ have invaded, any strategy introducing round-number dependence (whether on the responder or the investor's side) is, at best, strictly neutral.

(ii) If $ap_b > c$, then an investor who is reward-independent but round-number dependent (not shown) would be capable of invading. However, the strategy $\tilde{I}$, which is both reward-independent and round-number independent, can also invade while being strictly closer to the resident strategy (not shown). This means that round-number dependence is not parsimoniously evolvable in one step. Furthermore, once $\tilde{I}$ has invaded, the reasoning remains the same. At no point is round-number dependence parsimonious.

### E.5.2 Round-number dependence on the responder's side first

Here, the only possible scenario is that the strategy $I^*$, which invests in all games and all rounds in a reward-dependent manner, first invades on the investor's side. Once this occurs, there is selection pressure on the responder to reward at least some rounds of the trust game. A round-dependent responder $R^*_A$, who rewards only *certain* rounds of the trust game, is therefore capable of invading. However, in this situation, a responder $R^*$ who rewards *all* rounds of the trust game can also invade while being strictly closer to the resident strategy $R_0$ (see table SI.5). And once it has invaded, the resulting strategy profile is $S^*$, which is an equilibrium. Hence, $R^*_A$ is not parsimoniously evolvable.

In conclusion, when $A \notin \{\emptyset, Z^+\}$, there is no evolutionary path from $S_0 = (I_0, R_0)$ to $S^*_A = (I^*_A, R^*_A)$ through a sequence of parsimonious adaptive mutations. Hence, $S^*_A$ is not evolutionarily parsimonious from $S_0$.

## E.6 Parsimony is incompatible with subgame perfection

The parsimonious cooperative equilibrium profile, $S^*$, is not subgame perfect. This is due to the highly punitive nature of the investor strategy $I^*$. Any instance of a missing reward in a game leads to the permanent cessation of investment in that game. In an out-of-equilibrium situation where a responder deviates from $R^*$ by playing *Defect* once in a trust game, the strategy $I^*$ dictates that the investor should then play *Decline* in all future rounds of the game. Yet, the best response in this situation would be to continue playing *Invest*, since the responder will return to the $R^*$ profile and reward all future investments.

A more complex strategy on the responder's side, $R^{**}$, would be necessary to obtain a strategy profile $S^{**} = (I^*, R^{**})$ that is subgame perfect. To achieve this, $R^{**}$ must contain a self-punishing principle as follows (fig. SI.4):

- If the game is a by-product cooperation game or an inefficient cooperation game, play *Defect*.
- If the game is a trust game and the responder itself has always rewarded in the past, play *Reward*.
- If the game is a trust game, but the responder itself has defected at least once in the past, play *Defect* (self-punishing principle)

However, this subgame perfect profile $S^{**}$ is not parsimonious.

The strategy $R^{**}$ is characterized by its conditional response to both the nature of the game—it is game-dependent—and its own past behavior—it is also self-dependent. In contrast, $R^*$ is only game-dependent. Starting from an ancestral strategy that is not self-dependent, $R^{**}$ is always further away than $R^*$. For example, starting from $R_0$, which is neither game-dependent nor self-dependent, the framework of finite automata gives a distance of $d(R^{**}, R_0) = 4$, whereas for $R^*$, we had $d(R^*, R_0) = 3$.

Yet, self-dependence provides no advantage under any circumstances, even after random deviations of all sorts. The ability to punish oneself is at best neutral and at worst costly. As a consequence, in any evolutionary scenario, if a self-dependent strategy is capable of invading, there always exists a non-self-dependent strategy that can invade as well. Therefore, a self-dependent strategy can never be evolutionarily parsimonious. Parsimony is thus incompatible with subgame perfection.

# F Alternative definitions of parsimonious equilibria

In this section, we adopt a slightly different definition of parsimony, one that accounts for Maynard Smith's second condition—the interactions between rare mutants. We then explore its implications in two games.

Finally, we note that while a broader definition of parsimony is possible, it is likely not desirable.

## F.1 Parsimony when considering selection from rare mutants

Consider a game defined by a finite set of $n$ players, a tree representing the sequence of decisions, a set of terminal nodes with associated payoffs for each player, and a partition of the non-terminal nodes into $n + 1$ subsets—one for each player and one for "Nature's moves", representing events beyond the players' control. A player's strategy is defined by a mapping from the set of their nodes to the set of available behavioral policies.

Here, we adopt the same three assumptions as in the main text.

First, following the standard assumption of vanishingly rare mutations, the evolution of strategies in this game is approximated as a trait substitution sequence. Second, the mutation graph is assumed complete, meaning any strategy can potentially arise as a mutation from any other. Third, an infinite population is assumed, allowing only substitutions to strictly advantageous strategies, as mutants with neutral or deleterious effects have zero fixation probability.

Under these assumptions, the evolutionary transition from strategy $S_1$ to strategy $S_2$ is said to be *feasible* if and only if the following condition holds:

$$\begin{cases} P(S_2, S_1) > P(S_1, S_1), \text{ or} \\ P(S_2, S_1) = P(S_1, S_1) \text{ and } P(S_2, S_2) > P(S_1, S_1), \end{cases}$$

This means that $S_2$ either performs strictly better in interactions with $S_1$ or is strictly neutral in interactions with $S_1$ but performs strictly better when interacting with another $S_2$.

For clarity, an indicator variable $\sigma_{S_1 \to S_2}$ is used in the following, set to 1 if the evolutionary transition from $S_1$ to $S_2$ is feasible, and 0 otherwise.

The evolutionary transition from strategy $S_1$ to strategy $S_2$ is then said to be *parsimonious* iff (i) $\sigma_{S_1 \to S_2}$ (i.e., the transition is feasible), and (ii) the following condition is met:

$$\forall S_3 \neq S_2 \text{ with } P(S_3, S_1) > P(S_1, S_1),\ d(S_3, S_1) \geq d(S_2, S_1) \tag{7}$$

In other words, $S_2$ is the closest strategy that can successfully invade $S_1$.

As in main text, a strategy $S^*$ is then defined as an *evolutionarily parsimonious equilibrium* from an ancestral strategy $S_0$ if and only if the two following conditions hold:

1. There exists at least one trait substitution sequence from $S_0$ to $S^*$ that includes only parsimonious transitions.
2. $\forall S \neq S^* P(S, S^*) < P(S^*, S^*)$, i.e., $S^*$ is an equilibrium.

In what follows, we apply this alternative definition to two games.

## F.2 Hawk-dove game

In this game, interactions between mutants play a particularly significant role. Here, even without perturbations, strong asymmetry can emerge because rare assymetric mutants can pay the role of a perturbation.

As in the case described in the main text, evolution first converges toward a symmetric strategy, where the resident plays Hawk with the threshold probability $\hat{p}$ in every state. Once this resident is established, any mutant remains strictly neutral against it. In particular, a maximally asymmetric mutant following a Bourgeois strategy—always playing Hawk in state $A$ and always Dove in state $B$—is also neutral against the resident. However, this mutant performs strictly better against itself beause it avoids all conflict costs. And there is no symmetric mutant able to achieve the same benefit. Hence the evolutionary transition towards asymmetry is parsimonious.

Here, asymmetry can emerge without any external perturbation because the rare mutant itself creates the very perturbation that ultimately gives it an advantage. However, this emergence of asymmetry is only possible because assymetry does provide a solution to a coordination problem that could not be solved with a symmetric strategy.

By contrast, in the divide-the-dollar game, no transition from symmetry to asymmetry is ever parsimonious, as asymmetry is not needed to resolve a coordination problem. The only parsimonious equilibrium remains the symmetric one, even under this alternative definition of parsimony.

## F.3 Repeated prisoner's dilemma

In the repeated prisoner's dilemma, if we adopt the main text definition of parsimony, where interactions between mutants are neglected, cooperation can never evolve from a non-cooperative ancestral state. This is because a mutant capable of reciprocal cooperation gains no advantage against the resident, who never cooperates. Any benefit it derives comes solely from interactions with other mutants. By neglecting these interactions, we rule out the evolution of reciprocity.

This is actually a good property of this definition of parcimony. As many have argued (56, 71–73) and as I have discussed elsewhere (53, 54, 74), the repeated prisoner's dilemma is not a satisfying ecological model for the evolution of reciprocal cooperation. In this game, in the initial absence of recciprocal cooperation, individuals are never really under selection to cooperate reciprocally. Instead, if reciprocal cooperation does appear suddenly, like a hopeful monster, then rare mutant-mutant interactions may allow it to spread. But such an event is highly improbable (33, 74).

For this reason, most empirically minded researchers in the evolution of cooperation believe that reciprocal cooperation is far more likely to have originated in ecological settings where individuals faced immediate selection pressure to cooperate conditionally (53–56, 73, 75**?** ). This idea is captured by the investment game.

Still, for the sake of theory, it is worth considering how parsimony applies to the repeated prisoner's dilemma.

To this aim, consider three types of strategies (see figure SI.5 for their representation as finite automata):

- AllD, which always defects. - Grim, which cooperates in the first round and continues to do so as long as its partner has always cooperated in the past, but permanently stops cooperating after a single deviation by its partner. - $\text{Grim}_A$, a broad family of strategies that, like Grim, cooperate in the first round and then continue cooperating in every round that belongs to an arbitrary subset $A$ of strictly positive integers, defecting in all other rounds. However, if the partner defects even once in a round belonging to $A$, the strategy permanently ceases cooperation.

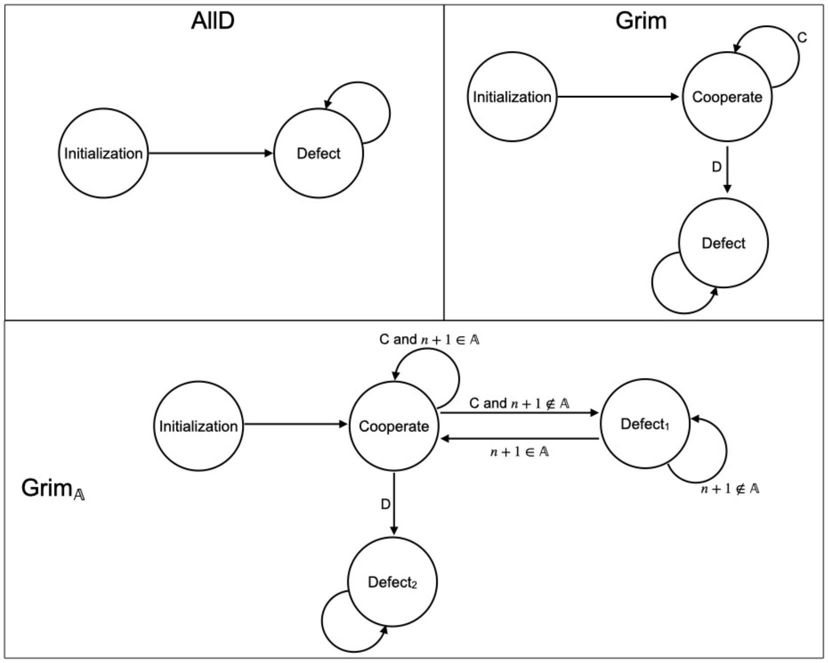

Figure SI.5: **Strategies in the repeated prisoner's dilemma represented as finite automata**. The distance between AllD and Grim is 4. The distance between AllD and $\text{Grim}_A$ is 7.

Here, we assume that $P(\text{Grim}, \text{AllD}) = P(\text{Grim}_A, \text{AllD}) = P(\text{AllD}, \text{AllD})$, i.e., we ignore the cost of cooperation in the first round against a pure defector. In this case, both types of reciprocal strategies, Grim and $\text{Grim}_A$, are neutral against a resident defector. This allows any strategy of the Grim family (plain or bizarre) to spread through the effect of mutant-mutant interactions, as each gains an advantage when interacting with itself (i.e., $P(\text{Grim}, \text{Grim}) > P(\text{AllD}, \text{Grim})$ and $P(\text{Grim}_A, \text{Grim}_A) > P(\text{AllD}, \text{Grim}_A)$). Thus, from AllD, evolution can lead either to plain reciprocity or to any one of the bizarre forms of reciprocity that condition behavior on the round number.

But parsimony eliminates all bizarre equilibria, leaving only plain reciprocity. The mutational distance from AllD to Grim is smaller than from AllD to $\text{Grim}_A$, as the latter requires the emergence of two new types of conditionality (conditionality to cooperation and conditionality to round number). The parsimony approximation thus assumes that the evolutionary transition towards bizarre forms of reciprocity will never occur. As a result, the only parsimonious equilibrium is plain reciprocity.

### F.4 Parsimony in finite populations: accounting for both occurrence probability and fixation probability

In principle, one could consider an even more general definition of parsimony.

For any possible mutation arising from a resident strategy fixed in a population, the probability of an evolutionary transition to that mutation—known as the probability of substitution—is the product of two probabilities:

(i) the probability that a single mutant carrying this mutation appears in the population, (ii) the probability that the mutation eventually fixes in the population after it has appeared as a single mutant.

In all previous analyses, the second probability has been neglected. The substitution probability of a strategy $S_1$ by a strategy $S_2$ has been assumed to depend solely on the probability of a mutant $S_2$ appearing in a population of $S_1$. This simplification follows from the assumption of deterministic selection, where a mutation fixates with probability 1 if it is advantageous and 0 otherwise.

However, this is a strong approximation, which in reality holds in no actual population. Even in very large populations, mutants are initially rare and can be lost due to genetic drift, and this probability of loss is not the same for all mutants. One may thus describe what the parsimony approximation would look like in the absence of this assumption.

For any single mutant present in a resident population, one can define its probability of fixation. These fixation probabilities can be computed under certain population models and approximations, but what matters here is that they can always be defined.

- If the mutant is strictly superior to the resident, this probability corresponds to the fixation probability of an advantageous mutation, which depends on the strength of its selective advantage. - If the mutant is neutral or disadvantageous against the resident but superior against itself, its fixation probability depends on its ability to cross an invasion barrier and subsequently fix, which in turn depends on both its costs and benefits. - If the mutant is strictly disadvantageous both against the resident and against itself, its fixation probability corresponds to that of a deleterious mutation.

Each possible mutation arising from a resident strategy is thus characterized by a substitution probability, which is the product of its mutation occurrence probability and its fixation probability once it has appeared.

Under this perspective, the parsimony approximation would then simply consist in assuming that whenever multiple substitutions are possible at a given point in the substitution sequence, the most probable substitution always occurs.

However, even though such an approach would have the advantage of generality, it would extend the parsimony approximation into a domain where its legitimacy is weaker, as it would also consider small quantitative differences in substitution probabilities. In this case, the approximation would become highly inaccurate and lose its heuristic value.

There would, in fact, be a paradox in constructing a fully detailed, minimally approximated model of substitution probabilities, only to then make the extremely crude approximation that assumes the most probable substitution always takes place.

This would misunderstand the role and heuristic value of the parsimony approximation. Its purpose is to eliminate strategies whose emergence is entirely implausible from a biological standpoint—strategies that are so complex that their occurrence probability is orders of magnitude lower than that of more biologically natural alternatives. In such cases, the role of fixation probability becomes negligible. This is why, in this article, I have retained the simplest possible definition of parsimony.