Smith, D. (2018). Prisoner's Dilemma and Cooperation. In T. Shackelford, & V. Weekes-Shackelford (Eds.), *Encyclopedia of Evolutionary Psychological Science* Springer.
https://doi.org/10.1007/978-3-319-16999-6_3757-1

Peer reviewed version

Link to published version (if available):
[10.1007/978-3-319-16999-6_3757-1](https://doi.org/10.1007/978-3-319-16999-6_3757-1)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the author accepted manuscript (AAM). The final published version (version of record) is available online via Springer Nature at https://doi.org/10.1007/978-3-319-16999-6_3757-1. Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

Prisoner's Dilemma and Cooperation

Daniel Smith

Bristol Medical School: Population Health Sciences, University of Bristol, Bristol, UK

Synonyms: Evolutionary game theory; Strategies for successful cooperation; Evolution of cooperation; Social dilemma.

Brief definition

The Prisoner's Dilemma (PD) is a two-player game where there is a conflict between individual and group interests. Overall, both players do better if they both cooperate, yet each individual player does better if they defect. However, if both players defect then they do worse than if they both cooperate, hence the dilemma. This deceptively simple game has been used extensively to explore the conditions under which cooperation can evolve.

Introduction

The widespread evidence for cooperation in the natural world is often seen as an evolutionary puzzle. This is because, all else being equal, it is fitness-enhancing to reap the rewards of others' cooperation without paying any of the costs. In a group of cooperators who freely help anyone else, selfish strategies possess increased fitness relative to cooperative strategies, and consequently spread through the population until only selfish types remain (West, Griffin, & Gardner, 2007). The

tragedy of this scenario is that all individuals would be better off (i.e., have greater absolute fitness) if everyone cooperated (Rand & Nowak, 2013). However, natural selection is inherently myopic and only concerned with immediate relative fitness differences between phenotypes, hence this rather pessimistic conclusion. Nice guys, it would appear, do indeed finish last.

However, cooperation is widespread throughout the biological world, from aphids and bacteria to yeast and zebra finches. Meerkats teach young group-mates to hunt and handle prey by deweaponizing otherwise dangerous scorpions; bees will aggressively defend threats to their nest even if they sacrifice their own lives in the process; and slime moulds, which are otherwise solitary, build communal structures when times are tough where only some individuals get the opportunity to spore and reproduce. Closer to home, our bodies are composed of highly-cooperative agents, with all our somatic cells facing no prospect of reproduction in order to pass on the germ-line. Even genomes and chromosomes can be thought of as collectives of cooperative individual genes.

Human beings are cooperative creatures *par excellence*, working together to build societies, engage in warfare (even at a high risk of mortality) and help strangers in need. We even extend this cooperation across the species barrier, by looking after pets who appear to give very little back in return. Even invisible deities or abstract ideological concepts frequently receive our cooperation. Given the ubiquity of cooperation, there must be some solution to the rather bleak conclusion that all life is destined towards selfishness. Without cooperation, the genes of our distant ancestors would not have collaborated, complex life would not have evolved, and the biological world would solely consist of short, microscopic, strands of DNA. Of course, we should not be misled into viewing the widespread cooperation in nature

through rose-tinted spectacles; exploitation is also rife in nature, with organisms frequently willing to manipulate, cheat or deceive others, and this aspect of behavior should not be overlooked. Given this, the widespread cooperation we observe in the biological world must have evolved because these cooperative strategies outcompeted their less-cooperative counterparts. Using the framework of the Prisoner's Dilemma, several such solutions will be discussed below.

First, the term 'cooperation' needs to be defined, as different authors often have different definitions of what constitutes 'cooperation' (and associated terms such as 'altruism'). Here, cooperation is defined as "a behaviour that provides a benefit to another individual (recipient), and the evolution of which has been dependent on its beneficial effect for the recipient" (West, Griffin, & Gardner, 2007, p. 662). This includes behaviors which are both 'mutually beneficial' (i.e., both the actor and recipient benefit from cooperation) and 'altruistic' (i.e., the recipient benefits but the actor pays a cost). These 'costs' and 'benefits' are measured in terms of individual fitness, rather than short-term gains and losses (e.g., resources), meaning that behavior which is costly in the short-term yet fitness-enhancing in the long-term (e.g., reciprocity; see below) would be classed as mutually beneficial, rather than altruistic.

This definition also includes the important proviso that said behavior must have 'evolved to benefit others', meaning that behavior which helps others as a by-product of otherwise self-interested behavior does not count as cooperation. For instance, imagine a scenario where a group of four organisms are being attacked; if none of the group retaliate, then the whole group perishes, while each individual who does retaliate increases their own probability of dying by 10%, while increasing the probability of group survival by 25% (Clutton-Brock, 2009). In this scenario the best strategy is for individuals to retaliate, regardless of the actions of others, as the costs

to retaliation (a 10% increase in mortality risk) are outweighed by the benefits (a 25% decrease in mortality risk). Thus, even though this retaliatory behavior may seem cooperative as it increases the survival of other group members, the benefits derived by others are an incidental by-product of otherwise self-interested behavior, so this should not be categorized as 'cooperation'.

The Prisoner's Dilemma

The Prisoner's Dilemma (hereafter PD) is one of the most commonly-used tools to explore the evolution of cooperation. The PD is an outgrowth of 'game theory', a branch of mathematics originated by John von Neumann and Oskar Morgenstern designed to analyze the best strategies when interacting with others, given different conditions. In its simplest one-shot version, the PD formalizes the scenario outlined above, where cooperation should not evolve. There are two players in the PD, both of which have to decide whether to cooperate or to defect. On average, both players do better if they both cooperate, yet there is always the incentive to defect as the fitness pay-offs are greater. Yet if both defect, both players earn less than if they both cooperated. These situations are known as 'social dilemmas' (Rand & Nowak, 2013), which pit individual and group interests against one another. The PD is by far the most common of these social dilemmas, but others are also possible (see below).

The canonical example of the PD (and where it gets its name) concerns two prisoners in the following scenario. Two suspected criminals are brought in for questioning regarding a serious crime and taken to separate rooms. The police have enough evidence to convict them of a minor crime, but not enough for conviction of a

serious crime. Each prisoner is then offered a bargain and can either cooperate with the other prisoner (keep quiet) or defect (testify that the other prisoner committed the serious crime). The outcomes of these actions depend on the behavior of the other player. If both cooperate and keep quiet, then each gets charged with the minor offence and serves one year in jail. If one cooperates and the other defects, then the defector walks away free while the cooperator spends ten years in jail for the serious crime. While if they both defect then they both get five years in prison. These sentences (or pay-offs) are displayed in figure 1.

In this scenario, defection is the only rational decision. Even though mutual cooperation results in the highest average payoff for both prisoners (one year in jail), there is always the temptation to defect if your partner cooperates, as then you would not spend any time in prison. Cooperation is therefore a risky decision because if you cooperate but your partner defects, then you spend ten years in prison. Consequently, regardless of the behavior of your partner, defection is the only rational option: if your partner cooperates, you can do better by defecting; while if your partner defects, the best outcome is to also defect.

		Player A	
		Cooperate	Defect
Player B	Cooperate	-1 / -1	0 / -10
	Defect	-10 / 0	-5 / -5

Figure 1: A pay-off matrix for the Prisoner's Dilemma. The behaviors at the top and the side of the matrix reflect the behavioral decisions made by each player (either cooperate or defect). Each cell represents the number of years spent in prison. The pay-offs for player A are shown in the top right of each cell (above the diagonal), while the pay-offs for player B are shown in the bottom left (below the diagonal). For each prisoner, regardless of how their partner acts, the highest pay-off is to defect, even though mutually cooperation is superior to mutual defection.

For generality, these outcomes can be displayed in an abstract pay-off matrix (figure 2A). There are four possible outcomes: Reward (R), where both parties cooperate; Temptation (T), where there is the temptation to defect on a cooperating partner; Punishment (P), where both parties defect; and the Sucker's pay-off (S), where an individual cooperates but their partner defects. These outcomes can be ranked according to their pay-offs: $T > R > P > S$. This ordering of outcomes defines the PD scenario (other orderings of pay-offs are discussed below). In addition, for a PD scenario two Rewards must be greater than a Temptation plus a Sucker's pay-off ($2 \cdot R > T + S$), otherwise in iterated games repeated turn-taking of Temptation then Sucker's pay-offs would lead to greater pay-offs than repeated mutual cooperation. The 'classic' pay-offs for these outcomes are $T=5$, $R=3$, $P=1$ and $S=0$ (figure 2B). This formalizes the conclusion above that, regardless of the behavior of your partner, it is in your best interests to defect (as $T > R$ and $P > S$). As the same holds for the

other player, they should also defect, regardless of the behavior of the first partner. In this one-shot situation, defection is an 'evolutionary stable strategy' (ESS; Maynard Smith, 1982), in that no other strategy can invade a population of defectors.

		Player A	
		Cooperate	Defect
Player B	Cooperate	<div style="display: flex; justify-content: space-between; align-items: center;"> R R </div>	<div style="display: flex; justify-content: space-between; align-items: center;"> T S </div>
	Defect	<div style="display: flex; justify-content: space-between; align-items: center;"> S T </div>	<div style="display: flex; justify-content: space-between; align-items: center;"> P P </div>

		Player A	
		Cooperate	Defect
Player B	Cooperate	<div style="display: flex; justify-content: space-between; align-items: center;"> 3 3 </div>	<div style="display: flex; justify-content: space-between; align-items: center;"> 5 0 </div>
	Defect	<div style="display: flex; justify-content: space-between; align-items: center;"> 0 5 </div>	<div style="display: flex; justify-content: space-between; align-items: center;"> 1 1 </div>

Figure 2: Pay-off matrices of the Prisoner's Dilemma. The pay-offs for player A are shown in the top right of each cell (above the diagonal), while the pay-offs for player B are shown in the bottom left (below the diagonal). A) This matrix displays the abstract pay-off structure in a Prisoner's Dilemma. R=Reward (mutual cooperation); T=Temptation (defecting on a cooperative partner); S=Sucker's pay-off (cooperating on a partner's defection) and; P=Punishment (mutual defection). The ranking of these outcomes for a Prisoner's Dilemma (from best to worst in terms of pay-offs) are: $T > R > P > S$. B) This matrix displays the 'classic' pay-offs applied in a Prisoner's Dilemma. These pay-offs were used in Robert Axelrod's tournament.

This situation seems to fit many of our intuitions regarding the fragility of cooperation and the benefits of reaping the rewards of others' cooperation. Take food-sharing among hunter-gatherers; it is always better to be the one who receives food without paying the costs of foraging (in terms of time, energy or risk of injury), even if this reduces the overall amount of resources brought back into camp. Or take trench warfare during World War I; it is always better to attack first (defect), rather than abstain (cooperate), as the enemy might attack first, even if this increases the risk of mortality for both sides.

The PD can also be extended to an n -person game – the ‘tragedy of the commons’ – where the same conclusion holds. To use an agricultural example, if farmers have cattle which graze on a plot of land, it would be optimal in the long-run for individuals to only allow the cattle to graze as much as can be grown back (cooperate). This would mean that resources remain stable over time, sustaining the long-term future of the population. However, if one farmer purchases more cattle and allows them to graze on the land (defect), their short-term material pay-off will be higher than the other farmers. Therefore, other farmers must also obtain extra cattle in order to compete, causing the resources of the pasture to deplete past sustainable levels. Real-world illustrations of this ‘tragedy of the commons’ abound: over-exploitation of natural resources, man-made climate change and tax avoidance are all obvious examples where short-term individual gain can have damaging long-term population-level consequences. As can be seen from these handful of examples, the Prisoner’s Dilemma can be an incredibly useful tool for modelling myriad different social dilemmas which pit individual and group interests against one another.

To summarize, in one-shot games defection is always the best strategy, which is why the evolution of cooperation is often seen as a puzzle. However, given that cooperation is widespread in the natural world, especially among humans, there must be solutions to this rather pessimistic conclusion. In particular, note that the main conclusion of the PD only holds under rather special circumstances: i) partners will never meet again; ii) partners are not related; and iii) partners interact randomly with others in the population (i.e., individuals cannot choose their partners). If these conditions are not met, then cooperation can evolve, despite the short-term benefits of defection. Several such solutions will be discussed below.

Solution 1 – Repeated interactions

In a PD scenario, defection is the only optimal strategy if interactions are not repeated (and players are not related and cannot choose their partners; see solutions 2 and 3, below). Clearly, however, in iterated games the long-term pay-offs to cooperation are greater than the short-term benefits of defection. Imagine a PD with the pay-off structure in figure 2B, repeated for 10 rounds with the same partners. Now consider two sequences, X and Y. In sequence X, Player A cooperates first, while Player B defects; player A therefore receives 0 points, while player B receives 5. In sequence Y, both players cooperate, so both receive 3 points. So far, it appears that defectors have greater fitness. However, in sequence X Player A now defects (as does player B), so from rounds 2 to 10 each player only receives 1 point per round. For sequence Y, in contrast, both players continually cooperate, so earn 3 points per round. Both players in sequence Y earn 30 points, while in sequence X player A receives a total of 9 points, while player B receives a total of 14. Therefore, if the PD is iterated, cooperators can have higher fitness than defectors.

However, in a *finitely* repeated PD there is a problem. If players know that the game will only last 10 rounds, then they have an incentive to defect on the last round, thereby increasing their overall pay-offs without harming their future interactions with the other player (because there are none). The other player of course also realizes this, so both players are likely to defect on the tenth round. This effectively means that the ninth round of the game becomes the 'final' round, as players will expect their partner to defect on the tenth round, which encourages both players to defect in this preceding round. This vicious cycle repeats back until the first round, in which

case the game simply returns to a one-shot PD and the only rational choice is to defect (Axelrod, 1984).

Repeated interactions may not therefore solve the problem of cooperation unless the game is repeated for an *unknown* length of time (although note that experimental work has shown that even in finitely repeated games levels of cooperation are sustained for longer than theoretically expected, although defection does increase as the known end-point approaches; Selten & Stoecker, 1986). If players do not know when the final round will be then the above issue does not arise. In this situation, players do not know with certainty when the final round will be, removing the temptation to defect on the final round, therefore increasing overall levels of cooperation as long as the probability of repeated interactions is sufficiently high. This corresponds to the common-sense notion that it is better to cooperate with someone if you plan to meet them repeatedly in the future, compared to someone you probably will not meet again.

A test of these predictions was formalized in Robert Axelrod's infamous tournament competition, in which various strategies were pitted against one another in an iterated PD game (Axelrod, 1984). There were two versions of this tournament. In the first version, 14 game theorists from fields of economics, psychology, political science, mathematics and sociology submitted strategies which were paired against one another for five games of exactly 200 moves each (with the pay-offs in figure 2B). After this first tournament, a second tournament was staged, which received 62 entries from six countries. In this second tournament each strategy was paired for five games of varying length, with an average of 151 moves per game. Example strategies submitted to these tournaments include:

- RANDOM: This program randomly cooperates and defects with equal frequency.
- GRIM: This strategy cooperates until the other player defects, after which GRIM is unforgiving and always defects until the end of the game.
- TFT (tit-for-tat): This program simply begins with 'cooperate', then repeats the previous move of its partner (i.e., if its partner defects, TFT will defect on the next round, while if its partner cooperates, so will TFT in the next round).
- JOSS: Similar to TFT, but instead of always cooperating when the other player cooperates, 10% of the time JOSS will defect after its partner's cooperation.
- TFTT (tit-for-two-tats): This strategy is equivalent to TFT, but rather than defecting after every partner's defection, it will only defect once its partner has defected twice in succession.
- TRANQUILIZER: This rule attempts to build up mutually beneficial cooperative relationships with its partner, while also defecting if its partner defects too often. If a cooperative relationship is established, however, TRANQUILIZER attempts an occasional defection to try and exploit its partner, hoping that the partner will forgive the infrequent transgression. If the partner does not retaliate these defections, then the frequency of defections increases.

In both of these tournaments, rather surprisingly the simplest program came out on top: tit-for-tat (TFT). This strategy received the highest overall score in both tournaments and was shown to be robust in additional analyses against different compositions of other strategies. Additionally, in an iterated evolutionary simulation,

where the strategies with the highest score in each generation produced more offspring in the next generation, TFT again out-performed all other strategies and left the most descendants. That is, TFT did well against all manner of opponents, under a wide range of parameters. TFT was able to succeed against other cooperative strategies by reciprocating in kind, while at the same time was effective against 'nasty' strategies as it could not be exploited easily (any defections were met by an immediate defection by TFT). More complex strategies which tried to exploit others were less successful as their defections were likely to set off chains of mutual punishment and defection.

In this iterated PD scenario TFT was found to be an ESS, in that it is able to resist invasion from any other strategy, provided the PD is repeated enough times. This is because no other strategy can receive a greater pay-off – and therefore greater fitness – than TFT in an iterated PD (for additional details and formal proofs, see Axelrod, 1984 and Maynard Smith, 1982). The pay-off of TFT against itself is repeated mutual cooperation. The pay-off of TFT against any other cooperative strategy will also be sustained mutual cooperation (i.e., CCCC), so this strategy cannot have higher fitness than TFT, so cannot invade. The pay-off of TFT against an alternating strategy (e.g., DCDC) will always be lower than mutual cooperation (as $T+S < 2*R$), so this strategy cannot invade TFT. Similarly, a strategy of mutual defection (i.e., DDDD) will also have a lower pay-off than mutual cooperation, so also cannot invade. However, it is important to note that other cooperative strategies can do equally as well as TFT, and this point will be returned to below.

Given the success of this simple TFT strategy, from this tournament Axelrod devised four 'guidelines' characteristic of successful strategies which are likely to be important in order to promote cooperation:

- Be nice: Never be the first to defect. Of the 15 highest-ranked strategies, all but one was 'nice', in that they were never the first to defect (the highest-ranked 'nasty' strategy finished eighth). Of the bottom-ranking 15 strategies, all were nasty.
- Reciprocate both cooperation and defection: That is, be both retaliatory and forgiving. It is important to be retaliatory so that other strategies do not exploit you. TFT does this well by punishing any defectors immediately. At the same time, it pays to be forgiving. If another player defects but then tries to cooperate afterwards, it is better to forgive them and cooperate in the future rather than risk repeated mutual defection. It is this quality that enabled TFT to perform better in the tournament than GRIM (which never forgave an initial defection by its partner).
- Don't be envious: Successful cooperation is not about obtaining a higher score than your opponent. A striking (although obvious, given some thought) fact is that TFT never beat any of its opponents; the best it could do was to draw (i.e., mutual cooperation), but in the process both partners could accrue a high score. Strategies which tried to beat their opponents were more likely to end up in cycles of defection and consequently earn lower scores.
- Be clear: Don't be too clever. TFT is a simple strategy to understand, while more complex programs are difficult to figure out and therefore predict their behavior (such as TRANQUILIZER). For cooperation to be sustained, clarity is key.

This tournament highlights that importance of reciprocity and repeated interactions for the evolution of cooperation, supporting the conclusions presented a decade

earlier by Robert Trivers (1971). However, despite the success of TFT it is not an ESS under all circumstances. Firstly, in a population of ALLD (always defect), a lone TFT cannot succeed as there are no other cooperative strategies to interact with. Therefore, assuming that the mutation rate is not too high and that cooperative strategies enter the population one at a time, defectors will always have higher fitness (Axelrod, 1984). As a concrete example, given five rounds of an iterated PD, a solitary TFT will only earn four points (TFT is a sucker once, then defects), while its ALLD partner will earn nine points. All other ALLD strategies in the population which only play against one another earn five points each. ALLD is therefore also an ESS, although TFT has a much wider basin of attraction once cooperative strategies have gained a foothold in the population. TFT therefore has trouble explaining the initial evolution of cooperation from a non-cooperative population, but it can explain its maintenance and spread once cooperation has crossed a certain threshold. However, from a population of defectors cooperation can evolve if two cooperative strategies emerge at the same time and preferentially cooperate with one another (cooperative assortment; see solution 3, below), or if the two players are related (kin selection; see solution 2, below).

Secondly, if there is 'noise' in the model, either in terms of mutation and drift or mistakes by other players (e.g., a cooperative strategy may accidentally defect) then TFT is not necessarily an ESS. In a population of TFT, the pay-offs for TFT and other cooperative strategies are equivalent, meaning that given some mutation rate other strategies can invade by genetic drift (Nowak & Sigmund, 1993). For instance, if ALLC (always cooperate) invades a population of TFT by drift, this will then trigger further evolutionary dynamics as 'nasty' strategies can exploit ALLC. Additionally, Axelrod's tournament is predicated on the assumption that strategies never make

mistakes. This is clearly not the case in the real-world as organisms may misinterpret or forget the behavior of their partner and subsequently defect when they should have cooperated. Indeed, after Axelrod's tournament it was found that TFT may not be the 'best' strategy, as PAVLOV (or, 'win-stay, lose-shift') was found to do better than TFT if mutation and occasional mistakes were included in the model. PAVLOV simply repeats its own previous behavior if it wins (R or T pay-offs), while shifts to the alternative strategy if it loses (P or S pay-offs). PAVLOV's success over TFT can be attributed to two factors. Firstly, it can exploit unconditional cooperators, while TFT only engages in mutual cooperation with these partners. This means that PAVLOV cannot be invaded by unconditionally cooperative strategies due to drift, unlike TFT. Secondly, PAVLOV can correct occasional mistakes, while for TFT any defection – even accidental – can result in repeated bouts of retaliation. There is therefore no single 'best' rule in an iterated PD, independent of the environment of other strategies, but rules such as TFT and PAVLOV do appear to be robust under a wide range of conditions.

Thirdly, these conclusions only hold in a simultaneous PD game, where both players make their decisions at the same time. While this situation may correspond to some types of cooperative behavior, other situations may reflect an 'alternating' PD, where individuals take it in turn to decide whether to cooperate or not (Nowak & Sigmund, 1994). For instance, bird parents often take turns obtaining food for their young, bouts of animal grooming are generally alternating, as are food-sharing situations in human and non-human animals. TFT may not be the best strategy in these scenarios, especially if reciprocation is not immediate and there are disparities in need. Need is important as the value of receiving cooperation is greater for those in need compared to those with an abundance of resources (Trivers, 1971). Take a

monetary example; \$1,000 is more valuable to a person in poverty than to a millionaire. Similarly, the costs to sharing are greater for individuals in need compared to individuals with lots of resources. Using the \$1,000 example again, the cost of sharing this amount is obviously greater to the person in poverty compared to the millionaire. Many PD scenarios assume that the pay-offs are equivalent between the two players, yet these value asymmetries greatly change the cooperative dynamics, making cooperation more likely to evolve (Trivers, 1971). For instance, in hunter-gatherer societies a small number of highly-productive individuals tend to provision others much more than they receive in return. Although this behavior appears altruistic, the costs to sharing these resources are low, relative to the benefits to the needy individuals receiving the food. Importantly, when the tables are turned and these productive cooperative individuals find themselves in need, due to illness or injury, they are more likely to receive resources than less-productive foragers (Gurven, Allen-Arave, Hill, & Hurtado, 2000). Similarly, among the Agta, a Filipino hunter-gatherer population, individuals in need were both less likely to share resources (D. Smith et al., 2016) and more likely to receive resources from others (D. Smith et al., 2018). Such need-based sharing may reflect reciprocal cooperation, but in circumstances beyond those modelled by simultaneous PD games such as Axelrod's tournament.

Nonetheless, despite these caveats regarding TFT, the basic conclusion that reciprocity and nice strategies can facilitate the evolution of cooperation has been repeatedly supported, both theoretically (see above) and empirically. As a few additional case studies: repeated economic games in a lab are associated with greater levels of cooperation (Rand & Nowak, 2013); stable hunter-gatherer camps with a high probability of repeated interactions are more cooperative than camps with

a high turnover in membership (D. Smith et al., 2016) and; truces in trench warfare in World War I were more probable when troops were stationed opposite one another for longer lengths of time (Axelrod, 1984).

Solution 2 – Kin selection

The paradox of the PD can therefore be overcome by repeated reciprocal interactions. A further solution is through cooperating with those who you share genes with, meaning that by helping relatives you can indirectly pass on your own genes to future generations. Although evolutionary geneticists such as Sewall Wright and JBS Haldane had an intuitive understanding that shared ancestry was important in explaining cooperative behavior, this approach was formalized by William Hamilton (1964) with his theory of 'inclusive fitness', often referred to as 'kin selection'. Hamilton demonstrated that genes for altruism could propagate through a population if these altruistic acts were directed towards genetic relatives. Thus, even though these altruistic acts harm the individual's own fitness (known as their 'direct fitness'), they can still evolve if it increases their inclusive fitness by increasing the fitness of their relatives (an individual's 'indirect fitness').

The likelihood of an altruistic trait spreading through the population depends on three factors: i) the benefit to the individual receiving the altruistic act (b), in terms of lifetime reproductive success; ii) the cost to the actor to performing said altruistic act (c), in terms of lifetime reproductive success, and; iii) the coefficient of relatedness between the actor and the recipient (r), in terms of the proportion of shared genes (e.g., 0.5 for parents/offspring/full siblings, 0.25 for aunts/uncles/nieces/nephews, 0.125 for cousins, and so on). Altruism is therefore expected if the benefits to the

recipient of a cooperative act, as a function of relatedness, outweigh the cost of the action. This is known as Hamilton's Rule, as is given by the simple formula: $b \cdot r > c$. To take an abstract example, if the cost of an altruistic act is a one-unit reduction in reproductive success for the actor, but a three unit increase for the recipient, the actor is likely to cooperate if the recipient is a full sibling, but not a niece or nephew. This is because, for a full sibling, 3 (the benefit to the recipient) times 0.5 (relatedness between siblings) is 1.5, which is larger than 1 (the cost to the actor), so the benefits outweigh the costs. For a niece or nephew, however, the benefit is 3 times 0.25 (0.75), which is lower than the cost to the actor, so cooperation would not be expected to evolve in this instance.

Kin selection is a powerful explanation for the existence of cooperation in nature, with much seemingly altruistic behavior explicable in terms of increasing an organisms inclusive fitness, even if it damages their direct fitness (West et al., 2007). Organisms caring for their siblings, self-sacrifice in eusocial insects, somatic cells working in service for the germ-line, and numerous other examples highlight the importance of kin selection in explaining cooperative behavior. Although it can be difficult to dissociate indirect and direct fitness benefits in humans – that is, we often reciprocally cooperate with our relatives (Rand & Nowak, 2013) – kin selection appears to be a powerful force shaping human affairs: kin-based nepotism is rife in human societies, including inheritance rules and regal succession; fictive kinship terms, such as 'brother' or 'sister', are applied to non-kin as a sign of friendship; while food-sharing among foragers is often targeted towards kin, even when controlling for reciprocal sharing (D. Smith et al., 2016).

This kin selection approach can also be formalized in a PD situation. In traditional PD scenarios players are assumed to be unrelated to one another. This means that

there are no indirect fitness benefits to cooperating with your partner, as helping your partner does not increase your indirect fitness ($r=0$). In contrast, when playing against relatives these indirect fitness benefits need to be taken into consideration, as they can alter the subsequent pay-off structure (Maynard Smith, 1982). Given the traditional PD pay-offs presented in figure 2B ($T=5$; $R=3$; $P=1$; $S=0$), when adopting an inclusive fitness perspective the indirect fitness benefits accrued by your partner also need to be taken into consideration. Therefore, if both players cooperate (R), the inclusive fitness perspective needs to sum both the direct fitness benefits (a pay-off of 3) and the indirect fitness benefits (a pay-off of 3 multiplied by the relatedness between players), meaning that the inclusive fitness pay-off is: $R+(R*r)$. The same applies for mutual defection: $P+(P*r)$. For a Sucker's pay-off, the indirect effects of your partner's Temptation pay-off needs to be considered, so: $S+(T*r)$. While for a Temptation, the Sucker's pay-off of your partners needs to be accounted for, so: $T+(S*r)$.

This is summarized in figure 3, using an example coefficient of relatedness of 0.5. Under this situation, although exploitation (T) is still preferable to mutual cooperation (R), being the sucker (S) is now superior to mutual defection (P). As mutual defection is now the worst-case scenario for both players, individuals are incentivized to cooperate. Although technically this is not a PD scenario anymore (but rather a different social dilemma known as a 'snowdrift' type game, see below), this example demonstrates how relatedness can promote cooperation, even in one-shot scenarios. All of this will, of course, vary by the coefficient of relatedness and the pay-off matrix. Cooperation will be less likely among more distant kin or unrelated individuals, while under different PD pay-off matrices (e.g., $T=5$; $R=4$; $P=1$; $S=0$) mutual cooperation may be the best course of action if r is high enough.

Figure 3: The inclusion of indirect fitness benefits changes pay-offs of the Prisoner's Dilemma. The left matrix shows the formulas for constructing the Prisoner's Dilemma pay-offs between relatives. If both players cooperate, then the inclusive fitness pay-off is an additive combination of both an individual's direct fitness (R) and the other player's pay-off multiplied by the relatedness coefficient ($R*r$). Given a value of $r=0.5$, the indirect fitness benefit to mutual cooperation is 1.5 ($3*0.5$), which gives a total inclusive fitness pay-off of 4.5 ($3+1.5$). The right matrix displays the pay-offs for each of the Prisoner's Dilemma outcomes given a value of $r=0.5$, using the pay-offs in figure 2B. Among close kin, an inclusive fitness approach changes the Prisoner's Dilemma to a snowdrift scenario (see figure 5), where mutual defection is no longer an evolutionary stable strategy. For simplicity, only the pay-offs to Player A are displayed here.

Solution 3 – Cooperative assortment

The pessimistic conclusion of the PD also only holds in a well-mixed population, where cooperators cannot choose who they interact with. If, however, organisms can 'rig the game' and only interact and cooperate with other cooperative organisms, then cooperation can evolve much more readily. The previous two solutions above can also be thought of as solving the dilemma via similar assortative processes; repeated cooperative interactions with the same individual, rather than interacting randomly with the population (reciprocity; solution 1) and interacting with close kin, as opposed to non-relatives (kin selection; solution 2). At a fundamental level,

cooperation can evolve if there is positive assortment in cooperative phenotypes; in the words of William Hamilton “kinship should be considered just one way of getting positive regression of genotype in the recipient, and that it is this positive regression that is vitally necessary for altruism. ... it obviously makes no difference if altruists settle with altruists because they are related (perhaps never having parted from them) or because they recognize fellow altruists as such, or settle together because of some pleiotropic effect of the gene on habitat preference” (Hamilton, 1975, p. 337). Thus, all theories for the evolution of cooperation are predicated on the principle of assortment.

However, while reciprocity and kin selection are theories of assortativity, they are not theories of *cooperative* assortativity as they do not require organisms to cooperate with others based solely on their partner’s level of cooperativeness. As will be described below, theories of cooperative assortativity can help explain the otherwise seemingly inexplicable cooperation between non-relatives who may not interact again in the future. This type of cooperation is seen routinely in modern market-based economies, including giving to charity, volunteering, purchasing goods and academic peer review.

At its most basic level, cooperative assortment occurs when cooperative phenotypes preferentially interact with other cooperative phenotypes. This is easier said than done, however, as mechanisms to establish and maintain this cooperative assortment are required to keep defectors at bay. The concept of ‘partner choice’ is relevant here, especially when compared to so-called ‘partner control’ (Barclay, 2013). Partner choice occurs when individuals can choose who to cooperate with (or to ‘walk away’ from non-cooperators), while under situations of partner control individuals cannot avoid non-cooperators, but rather have to control the effects their

partner's selfish behavior. In a repeated PD scenario like Axelrod's tournament, partners cannot choose each other, so strategies which limit exploitative partners by conditional cooperation and punishment, such as TFT, are necessary; this is an example of a 'partner control' situation. On the other hand, if individuals could choose their partners, or walk away from them when they wished, this would constitute a 'partner choice' situation. Therefore, if individuals can avoid uncooperative others, rather than having to control them, cooperation can evolve more readily, even in one-shot interactions (Rand & Nowak, 2013; Roberts, 2015).

One method to promote this cooperative assortment is through the use of reputation. Indirect reciprocity is one such mechanism, which is based on the principle of 'helping those who help others' (Nowak & Sigmund, 1998). If individuals acquire a reputation for cooperation by being cooperative, and this reputation is known by others, those with a reputation for cooperation are more likely to be cooperated with, irrespective of a lack of previous encounters. Theoretical models have indicated that cooperation can evolve via this process (Nowak & Sigmund, 1998), especially when combined with mechanisms of partner choice to avoid uncooperative individuals (Roberts, 2015). Results from several empirical studies (reviewed in Milinski, 2016) have found results consistent with indirect reciprocity. For example, non-cooperative individuals are less likely to be cooperated with in the future, while individuals also appear to engage in 'reputation management'; if their behavior will be made public to others they are more likely to cooperate. Such behavior is particularly apparent in online marketplaces, such as eBay, where individuals who cultivate a reputation for outstanding customer service are more likely to attract new customers.

A related theory based on reputation and cooperative assortment is 'competitive altruism' (also known as 'reputation-based partner choice'; Sylwester & Roberts,

2013). According to this theory, individuals first signal their cooperativeness, after which these cooperative individuals are more likely to be chosen for mutually-beneficial cooperative ventures. For instance, a forager may share their food widely with the camp in order to subsequently be chosen as a partner to share food with or to be chosen as a cooperative hunting partner. Although reputation-based partner choice and indirect reciprocity make similar predictions, such that seemingly costly displays of cooperation will be rewarded by cooperation in future encounters, the mechanisms are distinct. In reputation-based partner choice, individuals preferentially interact with cooperative individuals for future mutually-beneficial cooperative interactions with said cooperator. In contrast, indirect reciprocity assumes that individuals help cooperative others solely to enhance their own cooperative reputation so that others will cooperate with them, irrespective of future mutually-beneficial interactions with the same partner. Nonetheless, both are theories of cooperative assortment which can explain how one-shot cooperation among non-kin can evolve.

Theories of cooperative assortment are likely to have great relevance in understanding cooperation in modern market-based economies where much cooperation is between unrelated individuals and interactions may not be repeated. Note, however, that these theories based on cooperative assortment may not necessarily predict patterns of cooperation in small-scale societies, where cues other than cooperativeness, such as relatedness or reciprocity, can be used to select partners. For instance, among Agta hunter-gatherers individuals were more likely to share with *less* cooperative camp-mates, potentially reflecting need-based sharing or the avoidance of being indebted to others. Instead, they were found to preferentially share with kin, reciprocal partners and those in need (D. Smith et al., 2018).

Solution 4 – Communication and coordination

Communication and coordination between players does not change the logic of the PD in one-shot games. However, simply giving strangers 30 minutes to interact before playing a one-shot PD is enough for participants to predict how cooperative their partner would be (Frank, Gilovich, & Regan, 1993). This simple result highlights the importance of communication and coordination in shaping cooperative behavior. This is especially likely to be a factor in repeated games as individuals need to know that their partner will also cooperate. Unlike the ‘free-rider problem’, in iterated games there is a ‘problem of coordination’, as the issue is not necessarily that of free-riders having increased fitness relative to cooperators, but rather the difficulty in coordinating behavior for mutual benefit. In these situations, cooperation may be the best strategy for all parties, yet cooperation may not occur due to a lack of common knowledge over how others will behave.

The iterated PD can be thought of as a coordination game; individuals do best in repeated cooperative interactions, but in order to reap this benefit they need to know that their partner will actually cooperate. These problems require ‘meta-knowledge’ to solve. Put another way, the important point is not whether an individual knows that cooperation is the best course of action, but whether said individual knows that their partner also knows that cooperation is the best course of action and will act accordingly. Communication is therefore necessary to solve these problems of coordination, and numerous studies have shown that simply allowing players to communicate can promote cooperation (reviewed in E. A. Smith, 2010).

These coordination games can be modelled in an abstract pay-off structure (figure 4), similar to the PD. In these scenarios, defection is not the optimal strategy, as the

highest pay-off in these games is achieved by partners coordinating their behavior (R). As a hypothetical example, imagine two hunter-gatherers have the option of either cooperating to hunt stag or foraging for berries by themselves (the equivalent of defecting in this scenario): If both individuals cooperate then they are likely to successfully hunt a stag, which is a big haul for both partners (a pay-off of 5 each); if both players decide to forage individually for berries, then both receive a lower amount of food (a pay-off of 1 each); however, if an individual decides to hunt stag and their partner decides to pick berries, then the berry-picker still receives some resources (a pay-off of 2), while the stag-hunter has no chance of a successful hunt by themselves, so receives a pay-off of 0. Therefore, in order to reap the rewards of cooperation individuals need to coordinate their behavior.

		Player A	
		Cooperate	Defect
Player B	Cooperate	5 / 5	0 / 2
	Defect	2 / 0	1 / 1

Figure 4: An example of a coordination game pay-off matrix. Here, in order to reap the most rewards players need to coordinate their behavior. The pay-offs displayed here reflect a 'stag-hunt' scenario (see text) where the player who cooperates does worse than the one who defects ($R > T > P > S$), although other coordination games exist where any non-matching behavior is inferior to mutual defection, in which case $R > P > T = S$. The pay-offs for player A are shown in the top right of each cell (above the diagonal), while the pay-offs for player B are shown in the bottom left (below the diagonal).

One real-world example of such a coordination game involves foraging strategies among the Lamalera from Indonesia (Alvard & Nolin, 2002). In this society individuals can either fish individually or collectively hunt whale (which has a greater pay-off than solitary fishing). In order to coordinate behavior and ensure that cooperation (collective whale hunting) is profitable for all involved, complex sets of social norms exist among the Lamalerans, including specific rules over distributing whale meat. These social norms ensure that all Lamalerans are aware of the ‘rules of the game’, thereby coordinating their behavior for mutual benefit. While language obviously serves this purpose of coordinating social behavior (E. A. Smith, 2010), storytelling may be another such mechanism to broadcast this meta-knowledge, therefore promoting cooperation (D. Smith et al., 2017). A large proportion of stories told in hunter-gatherer societies appear to concern social behavior, particularly regarding extolling the virtues of prosocial behavior (such as cooperation and equality) and the punishment meted-out to norm violators (such as selfish or lazy individuals). These stories also appear to enhance cooperation, as among the Agta (a Filipino hunter-gatherer population mentioned earlier) higher levels of cooperation in an experimental game were associated with a greater proportion of skilled storytellers in camp. Thus, the importance of mechanisms to coordinate behavior for mutual benefit should not be overlooked when exploring the evolution of cooperation.

Multilevel selection and the evolution of cooperation

This chapter has focused on an “inclusive fitness” approach to understanding cooperative evolution (Hamilton, 1964; West et al., 2007). An alternative conceptualization to understanding the evolution of cooperation is through “multilevel

selection” (MLS; Hamilton, 1975; Sober & Wilson, 1998). MLS decomposes fitness into a “between-group” component and a “within-group” component (if the population is not structured into groups, then the between-group component is absent).

Selection favors selfish behavior within groups (because altruists have lower fitness than selfish types within groups), while selection favors altruistic behavior between groups (because groups with altruists have greater fitness than groups with fewer or no altruists). If between-group selection outweighs within-group selection, then seemingly altruistic behavior can spread through the population.

However, there are two important points to note here. Firstly, MLS models are mathematically equivalent to inclusive fitness models; the only difference is how they decompose fitness. MLS decomposes fitness into within-group and between-group components, while inclusive fitness decomposes fitness into direct and indirect fitness components. From an inclusive fitness perspective, the consequences of between-group interactions are part of an individual’s inclusive fitness. Secondly, the MLS definition of altruism is different to the inclusive fitness definition of altruism.

Altruism in an MLS framework is a behavior which lowers an individual’s fitness within a group, regardless of the between-group benefits to cooperation (Sober and Wilson 1998). Altruism from an inclusive fitness perspective, as described above, is a behavior which decreases an individual’s direct fitness, so can only evolve if cooperation increases an individual’s indirect fitness (Hamilton, 1964; West et al., 2007). It is important not to confuse the two: behavior may be altruistic from an MLS perspective, but not altruistic from an inclusive fitness perspective.

An example may make these points a little less abstract. Take between-group conflict. Imagine two groups in conflict over a limited resource (land, say), and individuals in each group can either cooperate (help in the conflict at a cost to self,

relative to group-mates) or defect (not participate, so pay no cost, but reap the rewards of their group-mates' cooperation). From an MLS perspective, cooperative individuals obviously have lower fitness than selfish individuals within groups, so this behavior is altruistic (as defined by MLS). However, groups with more cooperators are more successful in between-group conflict, so cooperative behavior may spread in the population if the strength of between-group selection is great enough.

From an inclusive fitness perspective, however, if groups are composed of kin, then this cooperative behavior can be understood in terms of individuals maximizing their inclusive fitness by increasing their indirect fitness, even at a cost to their direct fitness. This behavior is therefore still altruistic, but this time from an inclusive fitness perspective. Alternatively, if groups are not composed of kin, then this cooperative behavior can still be understood from an inclusive fitness standpoint as individual fitness also depends on group fitness (i.e., individuals from cooperative groups have greater fitness than individuals from less cooperative groups). Cooperative individuals may therefore increase their direct fitness by being part of a cooperative and successful group, even if said group contains some defectors who have even greater fitness than themselves. This behavior would therefore be altruistic from an MLS perspective, but not altruistic from an inclusive fitness perspective.

Each of the processes described above as solutions to the PD – repeated interactions, kin selection, and cooperative assortativity – can be described in a multilevel, as opposed to an inclusive fitness, framework (Sober & Wilson, 1998). In each case, groups of cooperators propagate, even though individual cooperators have lower fitness than defectors within groups. However, it is essential to remember that MLS is not an alternative process to inclusive fitness, but rather is simply a different perspective on the evolutionary process and an alternative way of doing the

math. The perspective one chooses to adopt will depend on the specifics of the system being studied: if groups act predominantly as unified wholes (such as whole organisms or, potentially, eusocial insect colonies), then an MLS perspective may be fruitful; if populations are not structured, then an inclusive fitness approach will be more appropriate; for cases in-between – where populations are structured into groups, but these groups do not act as unified wholes (such as human groups) – then it may be a matter of personal preference which perspective to adopt.

Limitations of the PD for understanding cooperation

While the PD has undoubtedly been a valuable tool for understanding the evolution of cooperation, it does possess some limitations. Firstly, not all cooperative behavior can be modelled as a PD. As discussed above, many cooperative situations resemble coordination game scenarios (figure 4), which do not suffer from the free-rider problem but rather suffer from problems of coordination. An alternative scenario, although still a social dilemma, is known as the ‘snowdrift game’. Imagine a situation where two drivers are stuck behind a snowdrift; as with the Prisoner’s Dilemma, individuals can either cooperate (dig through the snowdrift) or defect (not dig). While joint cooperation would mean clearing the snow faster, defecting and letting the other dig would be a superior strategy as no energy would be expended. However, in this case being the ‘sucker’ (the individual who cooperates while the other defects) is superior to mutual defection, as they still benefit by clearing the snow and getting home. The pay-off structure for this scenario is therefore $T > R > S > P$ (figure 5). In the snowdrift game defection is not the optimal strategy in one-shot interactions, resulting in a mixed population of cooperators and defectors (Doebeli &

Hauert, 2005). Accordingly, experimental findings have indicated that levels of cooperation are greater in an iterated snowdrift game compared to an iterated PD game (Kümmerli et al., 2007). Note also that this game is formally identical to the 'Hawk-Dove' game (Maynard Smith, 1982) or to the common childhood game of 'chicken' (where the first to 'blink' – or 'chicken out' – is the cooperator).

When 'cooperating' in this snowdrift game situation the benefit to the partner may be incidental, meaning that behavior in these situations may not be strictly cooperative, but rather reflects self-interest (Clutton-Brock, 2009; West et al., 2007). That is, players cooperate (dig through the snowdrift) in order to get themselves home because that is preferable to being stuck behind the snowdrift; the benefit to the other player is wholly incidental. Although cooperation in this scenario benefits others, it did not evolve to do so as the effect on the other player is an unintended by-product. The snowdrift game appears to characterize several real-world circumstances, particularly regarding producer-scrouter dynamics. In these situations, some organisms produce resources, such as food, information or cultural knowledge, which is then appropriated – or 'scrounged' – by others in the group. To take a foraging example, while neither party may wish to forage and have food taken by others, the costs to not foraging are greater for some individuals. These individuals would therefore 'blink' first in this game of chicken. For instance, compare an individual with multiple dependent offspring against another individual with no dependents. In this scenario the second childless individual has less need to forage compared to the first as the costs to not foraging for the individual with multiple children are larger as their household is in greater need of resources. Consistent with this idea, men from Ifaluk atoll with more dependent offspring were more likely to fish compared to those with few dependents, as the costs to not foraging were

greater for those with many mouths to feed (Sosis, Feldstein, & Hill, 1998). These alternative pay-off structures have been under-researched compared to Prisoner's Dilemma scenarios, despite their seeming applicability to several kinds of social interactions among a number of taxa (Alvard & Nolin, 2002; Clutton-Brock, 2009; Doebeli & Hauert, 2005).

		Player A	
		Cooperate	Defect
Player B	Cooperate	<div>3</div> <div>3</div> <div>3</div>	<div>5</div> <div>1</div> <div>3</div>
	Defect	<div>1</div> <div>5</div> <div>3</div>	<div>0</div> <div>0</div> <div>3</div>

Figure 5: The pay-off matrix for a snowdrift scenario. As with the Prisoner's Dilemma (figure 2), the Temptation pay-off is superior to mutual cooperation (Reward). However, in contrast to the Prisoner's Dilemma, mutual defection (Punishment) is inferior to being the Sucker. The ranking of these outcomes for a snowdrift scenario (from best to worst, in terms of pay-offs) are: $T > R > S > P$. The pay-offs for player A are shown in the top right of each cell (above the diagonal), while the pay-offs for player B are shown in the bottom left (below the diagonal).

In addition to their theoretical utility, the PD and associated games, such as the Public Goods Game, Ultimatum Game and Dictator Game, have been used extensively in experimental settings to empirically assess behavior in social dilemma scenarios (Camerer, 2003). Although these experimental protocols offer conceptual simplicity and clarity, in practice there are many difficulties in interpreting behavior in these games. Foremost among these are concerns over the ecological validity and interpretation of such experimental results. Taking issues of external validity first, it is

often not clear how, or even whether, behavior in these games reflects real-world cooperation. In these experimental games initial levels of cooperation are generally quite high; in Public Goods Games players tend to share half of their endowment, while in Dictator Game situations (where individuals simply share some proportion of resources with an unknown recipient) individuals still often offer around 20% of the stake, despite no threat of punishment or future interactions (Camerer, 2003). However, in the real-world such cooperative behavior is generally not observed; a recent field study based in Las Vegas found that individuals given a windfall of resources (casino chips) at a bus stop shared none of the chips with another individual waiting nearby, even after being instructed that they could share them with the other person if they wanted (Winking & Mizer, 2013). Additionally, there are difficulties concerning whether behavior in these experimental games reflects a single, stable, cooperative construct. For instance, it is not clear whether cooperation is a unitary phenomenon – and therefore stable from one context to the next – or whether it is more context-dependent – in which case experimental cooperation may not reflect real-world cooperation or cooperation in a different experimental game. Although some studies have attempted to overcome these criticisms, by demonstrating that game behavior predicts real-world cooperation and that cooperation in one context predicts cooperation in a different context (e.g., Peysakhovich, Nowak, & Rand, 2014), interpretation over the validity of these games is an area of active research and debate.

Conclusion

The PD has been widely used as a theoretical and empirical tool to explore the circumstances under which cooperation is likely to evolve. As of August 2018, a Scopus search for the term 'Prisoner's Dilemma' in titles, abstracts and keywords returned over 5,000 articles. It is likely that this fascination with this seemingly simple yet beguiling game will continue for years to come, offering insights into why humans and other animals both cooperate extensively with each other yet simultaneously aim to exploit one another when possible. The PD is an intuitive and fundamental insight into social behavior and the conflict between selfish and cooperative strategies; the dilemma is unlikely to be solved any time soon.

Cross-References: Evolution of Cooperation; The Prisoner's Dilemma; Robert Axelrod's (1984) *The Evolution of Cooperation*; Reciprocal Altruism and Cooperation for Mutual Benefit; Reciprocal Altruism; Evolution of Reciprocal Altruism; Strategies for Successful Cooperation; Reputation and Altruism; Game Theory; Iterated Prisoner's Dilemma Model; Kin Selection Hypothesis.

References

- Alvard, M. S., & Nolin, D. A. (2002). Rousseau's Whale Hunt? *Current Anthropology*, 43(4), 533–559. <https://doi.org/10.1086/341653>
- Axelrod, R. (1984). *The Evolution of Cooperation*. New York: Basic Books.
- Barclay, P. (2013). Strategies for cooperation in biological markets, especially for humans. *Evolution and Human Behavior*, 34(3), 164–175. <https://doi.org/10.1016/j.evolhumbehav.2013.02.002>
- Camerer, C. F. (2003). *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press.
- Clutton-Brock, T. (2009). Cooperation between non-kin in animal societies. *Nature*, 462(7269), 51–57. <https://doi.org/10.1038/nature08366>
- Doebeli, M., & Hauert, C. (2005). Models of cooperation based on the Prisoner's Dilemma and the Snowdrift game. *Ecology Letters*, 8(7), 748–766. <https://doi.org/10.1111/j.1461-0248.2005.00773.x>
- Frank, R. H., Gilovich, T., & Regan, D. T. (1993). The evolution of one-shot cooperation: An experiment. *Ethology and Sociobiology*, 14(4), 247–256. [https://doi.org/10.1016/0162-3095\(93\)90020-I](https://doi.org/10.1016/0162-3095(93)90020-I)
- Gurven, M., Allen-Arave, W., Hill, K., & Hurtado, A. M. (2000). "It's a Wonderful Life": Signaling Generosity among the Ache of Paraguay. *Evolution and Human Behavior*, 21, 263–282.
- Hamilton, W. D. (1964). The Genetical Evolution of Social Behavior. *Journal of Theoretical Biology*, 7, 17–52. [https://doi.org/10.1016/0022-5193\(64\)90039-6](https://doi.org/10.1016/0022-5193(64)90039-6)

- Hamilton, W. D. (1975). Innate social aptitudes of man: An approach from evolutionary genetics. In R. Fox (Ed.), *Biosocial Anthropology* (pp. 133–153). London: Malaby Press.
- Kümmerli, R., Colliard, C., Fiechter, N., Petitpierre, B., Russier, F., & Keller, L. (2007). Human cooperation in social dilemmas: Comparing the Snowdrift game with the Prisoner's Dilemma. *Proceedings of the Royal Society B: Biological Sciences*, 274(1628), 2965–2970. <https://doi.org/10.1098/rspb.2007.0793>
- Maynard Smith, J. (1982). *Evolution and the Theory of Games*. Cambridge: Cambridge University Press.
- Milinski, M. (2016). Reputation, a universal currency for human social interactions. *Philosophical Transactions of the Royal Society B, Biological Sciences*, 371, 20150100. <https://doi.org/10.1098/rstb.2015.0100>
- Nowak, M., & Sigmund, K. (1993). A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner's Dilemma game. *Nature*, 364, 56–58.
- Nowak, M., & Sigmund, K. (1994). The alternating prisoner's dilemma. *Journal of Theoretical Biology*. <https://doi.org/10.1006/jtbi.1994.1101>
- Nowak, M., & Sigmund, K. (1998). Evolution of Indirect Reciprocity by Image Scoring. *Nature*, 393(June), 573–577.
- Peysakhovich, A., Nowak, M. A., & Rand, D. G. (2014). Humans Display a “Cooperative Phenotype” that is Domain General and Temporally Stable. *Nature Communications*, 5, 1–8. <https://doi.org/10.1038/ncomms5939>
- Rand, D. G., & Nowak, M. A. (2013). Human cooperation. *Trends in Cognitive Sciences*, 17(8), 413–425. <https://doi.org/10.1016/j.tics.2013.06.003>

- Roberts, G. (2015). Partner Choice Drives the Evolution of Cooperation via Indirect Reciprocity. *Plos One*, 10(6), e0129442.
<https://doi.org/10.1371/journal.pone.0129442>
- Selten, R., & Stoecker, R. (1986). End behavior in sequences of finite Prisoner's Dilemma supergames A learning theory approach. *Journal of Economic Behavior and Organization*, 7(1), 47–70. [https://doi.org/10.1016/0167-2681\(86\)90021-1](https://doi.org/10.1016/0167-2681(86)90021-1)
- Smith, D., Dyble, M., Major, K., Page, A. E., Chaudhary, N., Salali, G. D., ... Mace, R. (2018). A Friend in need is a friend indeed: Need-based sharing, rather than cooperative assortment, predicts experimental resource transfers among Agta hunter-gatherers. *Evolution and Human Behavior*.
<https://doi.org/10.1016/J.EVOLHUMBEHAV.2018.08.004>
- Smith, D., Dyble, M., Thompson, J., Major, K., Page, A. E., Chaudhary, N., ... Mace, R. (2016). Camp stability predicts patterns of hunter-gatherer cooperation. *Royal Society Open Science*, 3, 160131.
- Smith, D., Schlaepfer, P., Major, K., Dyble, M., Page, A. E., Thompson, J., ... Migliano, A. B. (2017). Cooperation and the evolution of hunter-gatherer storytelling. *Nature Communications*, 8(1). <https://doi.org/10.1038/s41467-017-02036-8>
- Smith, E. A. (2010). Communication and collective action: Language and the evolution of human cooperation. *Evolution and Human Behavior*, 31(4), 231–245. <https://doi.org/10.1016/j.evolhumbehav.2010.03.001>
- Sober, E., & Wilson, D. S. (1998). *Unto Others: The Evolution and Psychology of Unselfish Behavior*. Cambridge, MA: Harvard University Press.

- Sosis, R., Feldstein, S., & Hill, K. (1998). Bargaining theory and cooperative fishing participation on Ifaluk atoll. *Human Nature*, 9(2), 163–203.
<https://doi.org/10.1007/s12110-998-1002-5>
- Sylwester, K., & Roberts, G. (2013). Reputation-based partner choice is an effective alternative to indirect reciprocity in solving social dilemmas. *Evolution and Human Behavior*, 34(3), 201–206.
<https://doi.org/10.1016/j.evolhumbehav.2012.11.009>
- Trivers, R. L. (1971). The Evolution of Reciprocal Altruism. *The Quarterly Review of Biology*, 46(1), 35–57.
- West, S. A., Griffin, A. S., & Gardner, A. (2007). Evolutionary Explanations for Cooperation. *Current Biology*, 17(16), 661–672.
<https://doi.org/10.1016/j.cub.2007.06.004>
- Winking, J., & Mizer, N. (2013). Natural-field dictator game shows no altruistic giving. *Evolution and Human Behavior*, 34(4), 288–293.
<https://doi.org/10.1016/j.evolhumbehav.2013.04.002>