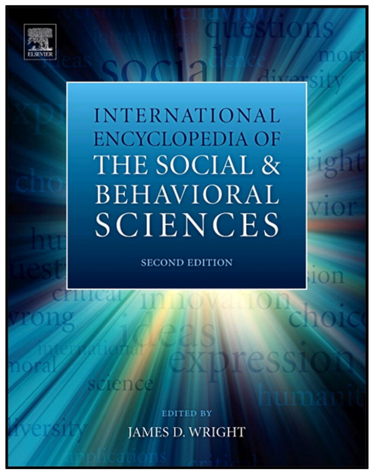

From Mathew, S., 2015. Human Cooperation, Evolution of. In: James D. Wright (editor-in-chief), International Encyclopedia of the Social & Behavioral Sciences, 2nd edition, Vol 11. Oxford: Elsevier. pp. 259–266.
ISBN: 9780080970868

Human Cooperation, Evolution of

Sarah Mathew, Arizona State University, Tempe, AZ, USA

Abstract

Human cooperation is a novel evolutionary puzzle because we cooperate with genetically unrelated individuals in groups that comprise millions of people. Direct reciprocity, especially when considering errors in behavior, has shed light on pairwise cooperation among well-known but genetically unrelated people. Three mechanisms have been identified to explain large-scale cooperation in humans. Cooperation can evolve by indirect reciprocity in groups as long as people can observe each other's behavior, or can garner honest information via gossip on who cooperates and defects in their interactions. Humans also have a disposition to cooperate and to punish those who do not, even at a cost to oneself, and such costly punishment can sustain cooperation in even larger groups of people who do not know much about each other. Cultural group selection explains the scale of human cooperation, why it is variable, and why norms enforced by sanctions are group-beneficial. Support for these theories has come from laboratory experiments using a variety of behavioral economic games, and from field studies in small-scale societies. Key open questions include understanding what characterizes goodness in indirect reciprocity, why gossip is sufficiently accurate, and why people are motivated to engage in costly punishment.

Introduction

Humans are much more cooperative than any other vertebrate species. Division of labor and delayed exchange of valuable commodities occurs in virtually every human society. The sick and injured are fed and cared for. Public capital facilities like communal shelters, ceremonial grounds, and defenses are commonplace in human societies. Such cooperation is virtually absent in other vertebrates. Humans even help strangers – donate blood, give directions, offer their seat, engage in philanthropy – more than most other vertebrates help individuals they have known their entire life. Language too is a low-cost signaling system, made possible because cooperation ensures that communication is sufficiently honest. Vertebrates that do not cooperate are limited to costly, self-verification of signals.

This high degree of cooperation makes humans an ultra-social species. Ultrasociality is rare in nature because of the detrimental effect of free riders – individuals who take the benefits of cooperation without paying the cost. Nonetheless when ultrasociality does occur, it has spectacular ecological consequences. Two eusocial taxa – ants and termites – make up one-third of the terrestrial animal biomass of the planet. Humans have higher biomass than any terrestrial mammal aside from the domestic cows we use in farming. Thus, if the free rider problem is solved, there are tremendous evolutionary gains to be had from cooperation.

Human ultrasociality is a novel evolutionary puzzle because it is not explained by genetic relatedness. Most ultrasociality in nature occurs among genetically related individuals (Smith and Szathmary, 2000). In multicellular organisms, cells are clones. In eusocial insects, workers tend to be siblings. But humans cooperate in groups comprised of hundreds, thousands, and even millions of genetically unrelated individuals. A growing body of research suggests that our extreme reliance on culture – socially transmitted knowledge, beliefs, and norms – might be the key ingredient that fueled this transition (Richerson and Boyd, 2005; Henrich and Henrich, 2007; Bowles and Gintis, 2011). Piecing together this puzzle has illuminated a diverse range of human phenomena such as our proclivity to punish wrong doers, our attention to fairness, laws and norms, morality and justice, and our penchant for gossip.

How Do We Know That Humans Are Cooperative?

Key insights on human cooperation have come from economic experiments, which have demonstrated that people all over the world violate the assumption that humans have purely self-interested preferences (J. Henrich et al., 2004; Camerer, 2003). In the Ultimatum Game, two anonymous players get to divide a sum of money. The proposer offers a division. If the responder accepts the offer, the two players get their proposed shares. If the responder rejects, they both receive nothing. If people are money-maximizers the lowest possible offer will be made and any offer will be accepted. Another game commonly used is the dictator game. Here the experimenter gives one of the players – the dictator – a sum of money. The dictator can either keep all of it for herself, or give some to the second player – the recipient. If players are money-maximizers, dictators will keep their full endowment. A third game is the public goods game. Here a group of participants are given a sum of money. They may contribute a portion of this endowment to a common pool, which is multiplied by the experimenter and divided equally among players regardless of their personal contribution. The multiplication factor is set so that a player can maximize her income by not contributing anything. In a seminal study Henrich et al. (Henrich et al., 2005) conducted these three experiments across 15 nonindustrialized societies including mobile hunter-gatherers, small-scale horticulturalists, acephalous pastoralists, and agriculturalists. People behaved quite differently in different societies, but in no society was their behavior consistent with money-maximization. Proposers offered more than the minimum possible amount, responders in many societies rejected low offers, dictators transferred money to the recipient, and participants contributed to the common pool.

International Encyclopedia of the Social & Behavioral Sciences, Second Edition, 2015, 259–266

Research on contemporary small-scale foragers also paints a picture of societies that are remarkably cooperative. Among the Shiwar, a forager-horticulturalist society in South America, more than 50% of adult men had disabilities that kept them from foraging for more than a month (Sugiyama, 2004). In any other vertebrate society, such a lengthy disability would mean death. Human foragers instead regularly survive these periods because others provide for them during such times. An analysis of the Hiwi and Ache, two forager societies in South America suggest that there is extensive food sharing beyond the household (Gurven, 2004). Among the Hiwi, in more than 75% of the foraging events, some amount of food was given to another household. In about 15% of foraging events, 90–100% of the food acquired is given to another household.

The human life history itself would be impossible without intensive cooperation, so much so that humans are considered to be cooperative breeders (Hrdy, 2009; Hill and Hurtado, 2009). Compared to the typical primate, humans invest in higher quality offspring. Human offspring take longer to reach sexual maturity and become self-sufficient foragers only by their late teens, but they have larger brains and longer lifespan. One would expect that such prolonged care of offspring would require longer inter-birth intervals – the usual pattern observed in mammals whereby quality comes at the expense of quantity. But surprisingly humans have higher fertility and shorter inter-birth intervals than chimpanzees, even though chimpanzee juveniles need no care. In a synthetic review of human life history traits, Kaplan et al. (2000) posit that food transfers are at the nexus of our escape from this trade-off. Post-reproductive women and adult males produce more food than they consume. Reproductive women and children produce less than they consume. Food transfers from these net-producers to these net-consumers allow reproductive aged women to support multiple dependent offspring at a time, and allow children to stay dependent for long.

Property rights are omnipresent in human societies of all types, from hunter-gatherers to nation states. There are tremendous benefits to be had if community members do not steal each other's unguarded property. People would not have to constantly guard or fight for their property and can utilize this time in production. Cooperation of this form must have been present among Paleolithic foragers. People would be disinclined to create complex tools and technology that require time and effort now but would pay off later, if they could lose these items to predatory band members. Human males spend far less effort staying in close proximity to their mates for mate-guarding than do males of most other animals. Free of the need to always mate-guard, men can spend their day foraging separately from women, specializing in different foods than women do.

Why Is Cooperation an Evolutionary Puzzle?

If cooperation is beneficial, and organisms that cooperate are immensely ecologically successful, why does evolution seldom favor cooperation? To see why, consider the payoff structures that organisms may encounter in nature that creates opportunities to benefit from cooperation. The payoff structure is the cost and benefit for each of the actors, given their behavior and the behavior of their partner(s). In some forms of cooperation, individuals have no incentive to defect (i.e., to not cooperate) as long as their partner cooperates. Table 1 illustrates such a payoff structure, also referred to as a mutualism or coordination game. Here two foragers have the option of either finding small game on their own, or coordinating with each other to hunt larger game. Assume that if they attempt to hunt the larger prey on their own they will fail. As long as the two foragers can communicate that they are willing to coordinate their foraging attempt, they have no reason to hunt alone. Realizing the benefits of cooperation in this case is easy: mutually agree to cooperate.

Table 1 Mutualistic interaction

		Forager 1	
		Hunt alone	*Hunt together*
Forager 2	Hunt alone	7, 7	7, 0
	Hunt together	0, 7	10, 10

An illustration of mutualistic payoffs where as long as one player cooperates, it is in the best interest of the other player to cooperate.

Table 2 Prisoner's dilemma game

		Player 2 (acquired fish)	
Prisoner's dilemma payoff structure		*Keep fish*	*Share fish*
Player 1 (acquired tubers)	Keep tubers	7, 7	12, 5
	Share tubers	5, 12	10, 10

An illustration of payoffs in a prisoner's dilemma where it is in a player's best interest to defect, regardless of whether the other player cooperates or defects.

Problematically, in many of the opportunities for cooperation in nature, individuals have an incentive to defect even when their partner cooperates. Such payoff structures are best illustrated by the prisoner's dilemma game (Table 2). Here one forager has spent the day fishing and the other acquiring tubers. At the end of the day, they can eat what they have personally acquired or they can share a portion of their day's catch with each other. Assuming that their optimal daily diet is a mix of tuber and fish, their joint payoff is maximized if they both share. But, forager 1 is better off keeping her fish if her partner keeps the tubers, but also if her partner shares the tubers. Likewise for forager 1. Thus both foragers will keep their catch forgoing the big gains they as a pair could have had from cooperating. The *n*-person public goods game illustrates this tension for a group of people: the group maximizes its payoff if everyone contributes to the public good, but each person has a higher payoff if they take the benefit without paying the cost of producing the public good.

Cooperation in a prisoner's dilemma or public goods game presents an evolutionary puzzle. To obtain the benefits of cooperation individuals have to behave in a way that maximizes the pair's or group's payoff, not their own. Evolutionary pathways to cooperation identify mechanisms by which cooperators can end up with higher payoffs (and consequently higher fitness) than defectors despite this disadvantage.

The Evolution of Small-Scale Cooperation

Kin Selection

Kin selection (Hamilton, 1964) explains a wide range of cooperation that occurs in nature, and plays an important role in human cooperation too. Altruistic behavior can evolve if individuals preferentially direct their behavior toward kin. This is because if a person carries the cooperative allele, her relatives are more likely than a random individual in the population to also possess the cooperative allele. This means that if individuals cooperate only with kin, cooperators will be more likely than defectors to also receive the benefits of cooperation. Such positive assortment allows cooperation to evolve.

Although nepotism is one aspect of human cooperation, it does not fully explain human cooperation. Nepotistic sentiments do exist. The meaning of the word – handing favors unfairly to relatives – illustrates the struggle to suppress this form of cooperation in order to organize cooperation at a larger scale. Food sharing in hunter-gatherers suggests that humans have evolved other pathways to cooperation that may often compete with our nepotistic motivations. A study among reservation-dwelling Ache of Paraguay (Gurven et al., 2001) shows that offspring, siblings and parents get a larger percentage of an individual's share of a resource item. However, close kin outside of the nuclear family like nieces, nephews, aunts, uncles and cousins receive less than what an unrelated band member receives. Another study conducted in the same population (Allen-Arave et al., 2008) show that foragers do not preferentially share with households that have their kin. Instead, foragers give away food to households that need food, and to households that have given them food in the past. Moreover, when there are disadvantages to cooperating with kin, humans are able to manage even high-stakes cooperation with nonkin. Among the Turkana, a nomadic pastoral society in East Africa, warriors venture on stealth missions with two to five other men, in which they plan to take cattle of neighboring ethnic groups. These missions can be risky, and the team needs to engage in a firefighting if detected. Although a high degree of cooperation and trust in one's companions is required, such teams are often comprised of age-mates and friends, not brothers. When I have enquired about this to Turkana participants I interviewed, they explain that one of the brothers needs to look after the herds, and that if the raiding party happen to be killed on this mission, a loss of two sons would be catastrophic for a family.

Misfiring of the nepotistic psychology in novel environments also does not satisfactorily explain human cooperation. Based on the high degree of cooperation in foraging societies, evolutionary social scientists have often assumed that humans evolved in foraging bands comprised of kin (e.g., Maner and Kenrick, 2010). If humans had evolved in kin-based bands, then the misfiring (Barkow et al., 1992) of the nepotistic psychology in present-day environments can account for human ultrasociality. However, this is an unlikely scenario. First, there is no evidence that Paleolithic foragers lived in kin-based bands. Data from two contemporary hunter-gather societies, the Ju/Hoansi and Ache, show that only up to a quarter of the band is comprised of a person's kin (Hill et al., 2011). The other half of the band is affines, and the remaining quarter is neither kin nor affines. Second, in acephalous nomadic pastoral societies like the Turkana, multiple settlements cooperate in warfare (Mathew and Boyd, 2011). Turkana settlements disperse and regroup seasonally like hunter-gather bands. With a band-level cooperative psychology, we should expect the Turkana to cooperate on the scale of settlements. Yet, an average Turkana raid has warriors from four different settlements. Lastly, people in societies with higher market integration and larger community size are more generous in the dictator game and more willing to punish unfair offers in the ultimatum game when playing with anonymous community members (Henrich et al., 2010). The misfiring hypothesis would suggest that people living in small-scale societies that are more similar to ancestral social environments should cooperate in anonymous interactions with their peers at least as much as people in large market-based societies.

Direct Reciprocity

Direct reciprocity (Trivers, 1971) can sustain cooperation among pairs of individuals who have the opportunity to interact repeatedly. Direct reciprocity is best exemplified by the *Tit for Tat* (*TFT*) strategy, where an individual cooperates in the first round in a repeated prisoner's dilemma game, and afterward acts just as her partner acted in the previous round (Axelrod and Hamilton, 1981). Reciprocal strategies like *TFT* can outcompete defectors because its practitioners are more likely to accrue the long-run benefits of repeated cooperation than defectors who attempt to exploit them in the first round but lose out afterward. Although there are limited examples of reciprocity in other animals, some forms of human cooperation fit with reciprocity. For instance, among the Hiwi and Ache, the amount of one family's food production given to another family correlates with the amount of the second family's food production given to the first family, controlling for biological kinship, geographic proximity, and the overall amount of this resource the family gives away (Gurven, 2004). Consistent with reciprocity models, humans have 'cheater detection' systems for social exchange – specialized cognitive systems to recognize when a partner has violated a social contract (Cosmides et al., 2010).

Although *TFT* was shown to be the winning strategy against a host of other strategies (Axelrod and Hamilton, 1981), reciprocity in nature does not follow its prediction too closely. The condition for *TFT* to be an evolutionarily stable strategy against a pure defector strategy is $wb > c$, where b is the benefit the recipient gets, c is the cost for the giver, and w is the probability the pair will interact again. According to this condition if a pair will interact twice on average, they could be as cooperative as siblings. If so, there should be rampant high-stakes pairwise cooperation in humans and other animals. Yet, reciprocity has rarely been documented in other animals, although repeated interactions occur in many animal societies (Clutton-Brock, 2009). Humans display conditional cooperation, whereby we are more likely to cooperate with partners who cooperate, but the strategy we utilize pays attention to much more than what occurred in the previous interaction. For instance, people pay attention to their partner's intent, not their behavior (Falk et al., 2008). Friendship is an appropriate domain for reciprocity – it is a long-term relation between two unrelated individuals characterized by mutual aid and companionship. Yet, a study conducted in USA, Japan, and China indicate that a characteristic of friendship is that people do not track each other carefully and tolerate imbalances in help received and given (Xue and Silk, 2012).

Why is there a mismatch between reciprocity theory and real-world pairwise cooperation? One reason is that in reciprocity models which strategy prevails depends on what mix of strategies are allowed to compete. This ambiguity prompted a recent study to examine what occurs when all possible strategies are allowed (van Veelen et al., 2012). The analysis by van Veelen et al. reveals that even when there are considerable gains to cooperation, there is only moderate cooperation in the population. The surprising pattern is caused by strategies that arise via mutation that can cooperate and thrive just as well as the common strategy, but once common, pave the way for less cooperative strategies. The population thus cycles between states that are dominated by more and less cooperative strategies. Interestingly, the authors find that if there is a small degree of assortment, repeated interactions yield much higher levels of cooperation. In humans, such assortment can arise if interactions are directed toward kin. Importantly, humans have yet another pathway for assorting. Because we acquire information through social learning, and learn what the local norms are about how to behave from our local group, interacting with cultural group members can also lead to assortment (Boyd and Richerson, 2010).

Another reason why real-world pairwise cooperation is unlike *Tit for Tat* is because *TFT* is a poor strategy when players make mistakes (Fudenberg and Maskin, 1990), fueling a quest for a variety of alternative strategies (reviewed in McElreath and Boyd, 2008; Nowak, 2006). Real-world interactions are plagued by genuine mistakes. People intend to cooperate, but may defect instead. For example, Sally agrees to pick Tom from the airport, but wrote the date incorrectly in her calendar and does not show up when Tom arrives. Such errors are called implementation errors. In another kind of error, people cooperate, but their partner may perceive it as a defection. For example, Sally arrives at the airport half an hour late because of traffic. Tom waited for 20 min and then took a taxi home. These are labeled perception errors. Resolving these errors have led to a diverse range of strategies that are not just retributive like *Tit for Tat*, but are additionally contrite – sorry for mistakes made and willing to make amends (Boyd, 1989), forgiving – occasionally cooperating with those who defect (Nowak and Sigmund, 1992), and flexible – shifting from cooperation to defection and vice versa based on remuneration under these tactics (Nowak and Sigmund, 1993). More investigation on this topic could illuminate further the emotions and behavioral strategies that underpin pairwise social relations in humans.

The Evolution of Large-Scale Cooperation

Human cooperation cannot be explained by nepotism and direct reciprocity alone. Humans cooperate with strangers, and cooperate in groups comprising millions of genetically unrelated individuals. In state societies, people regularly trade with anonymous strangers, and through taxes, contribute toward public goods like defense and capital facilities, and provide for individuals who are needy. Even without coercion from the state, people voluntarily cooperate: they donate blood, contribute to charity, assist in disaster relief, organize community events, protest unjust laws, give up their seat on a bus for an elderly person, return items to lost-and-found locations, give directions to travellers, etc.

Even politically uncentralized societies can organize cooperation at an impressively large scale. Among the Turkana, an acephalous tribal-scale pastoral society in East Africa, in addition to the clandestine stealth raids described earlier, men also mobilize for large-scale raids in which they engage in a firefight with neighboring ethnic groups (Mathew and Boyd, 2011). Such raids are a high-stakes form of cooperation – 1% of participants get killed in these raids, and 20% of all males born in the study population die in warfare. Although individuals take great risks in combat to overcome the opponent's resistance, all combatants can acquire the loot. Some gains like territory and deterrence are obtained even by noncombatants. Yet, the scale of Turkana raids is large. An average raiding party comprised of 315 warriors, who came from five different age groups, four different settlements, and three different subterritories of Turkana society. These men do not all know each other. Moreover, a warrior has on average only four close kin with him in such a raiding party.

Three mechanisms have been identified to explain the evolution of large-scale cooperation in humans: indirect reciprocity, punishment, and cultural group selection.

Indirect Reciprocity

Humans not only help those who helped them, but also help those who helped *others*. Alexander (Alexander, 1987) posited that such indirect reciprocity could explain society-wide cooperation enforced by moral systems that are based on how people treat others. Nowak and Sigmund (Nowak and Sigmund, 1998) theoretically demonstrated that indirect reciprocity can be evolutionarily stable. In this model, a pair of individuals is sampled at random from the social group. One of them is assigned to be the donor, and the other the recipient. The donor can cooperate by conferring a benefit b at cost c to the recipient, or the donor can defect. In the next round, another pair is sampled at random, and the game continues. Each act of cooperation increases a person's 'image score' while defection lowers this score. It turns out that a 'discriminator' strategy, which cooperates if its partner has a sufficiently high image score, can be evolutionarily stable against defectors if $qb > c$. Here q is the probability that the donor knows the image score of the recipient. Image-scoring was found to perform well only under a narrow set of conditions, and so has been replaced by 'standing' strategies which attend also to information on the donor's own score and to whether the recipient's past defections were justified (Panchanathan and Boyd, 2003; Leimar and Hammerstein, 2001). Nonetheless, the general consensus from the models is that cooperation can evolve by indirect reciprocity in communities where people can observe each other's behavior, or can garner honest information via gossip on who cooperates and defects in their interactions.

Key features of human cooperation are consistent with indirect reciprocity. Laboratory experiments have demonstrated that participants are motivated to help a novel partner who has helped other players. For instance, in one study (Ule et al., 2009) participants were paired at random with another participant each round. In one of the treatments the player in the donor role could help, do nothing, or harm one's partner. The majority of participants (57%) were other-regarding discriminators who helped if their partner had a reputation for helping others, and passed up on the opportunity to help if their partner did not have a history of helping others. A related observation is that humans are more cooperative when there are cues that they are being observed, suggesting that reputation matters. In an anonymous one-shot experimental study, participants were more generous to a stranger merely when stylized eye spots were drawn on the computer screen in which they took the study (Haley and Fessler, 2005). Nonverbal cues like mutual eye-gaze and touching increased contributions in a public-goods game (Kurzban, 2001). These subtle cues that we are being watched increase cooperation beyond the laboratory too (Bateson et al., 2006). Concern for reputation fits indirect reciprocity. However, individuals who strategically cooperate only when their own reputation is at stake resemble the self-interested discriminator strategy in the study by Ule et al. where 10% of players helped their partners when their own reputation was poor rather than when their partner had a good reputation. More theoretical work is required to understand the prevalence of self-interested discriminators and why they do not undermine cooperation sustained by other-regarding discriminators.

The above models of indirect reciprocity only explain pairwise helping among members of a social group. Can indirect reciprocity sustain large-scale cooperation? To answer this, one study examined players' behavior when rounds of the indirect reciprocity game alternated with rounds in which the group played a public goods game (Milinski et al., 2002). Players were much more cooperative in this game than if they were playing consecutive rounds of the public-goods games. Thus, the reputation consequence in the pairwise game was motivating players to behave well in the *n*-person game too. Panchanathan and Boyd (2004) theoretically demonstrated that indirect reciprocity can indeed support large-scale cooperation. If individuals in a social group engage in both *n*-person and pairwise cooperation, and if those who fail to contribute in the public-goods stage are denied help in pairwise exchange, then contributing to public goods is evolutionarily stable. Attitudes toward Turkana warriors who shirk on the battlefield are consistent with indirect reciprocity's role in sustaining *n*-person cooperation (Mathew and Boyd, 2014). In vignette studies, Turkana subjects were asked either about a fictitious coward who lagged behind others on a raid, or about an unskilled warrior who put in the effort but did not perform well. Participants were more likely to say that they would not help the coward if he needed to borrow an animal from them at a time of need. Being denied such help is consequential for a Turkana herder. Herdsmen lose their stock in raids, epidemics and droughts, and without aid from their relatives and community they cannot rebuild their herds. The threat of losing such assistance must motivate warriors to not shirk on the battlefield.

Indirect reciprocity has illuminated why humans mind other's business, and how this enables cooperation, but there remain three open questions. One question is why indirect reciprocity is even needed to sustain pairwise cooperation, when direct reciprocity could suffice. Second, it is clear that cooperation should be somehow contingent on the recipient having a good standing, but what is the best system of assigning 'goodness'? For instance, if I am in good standing and I am kind to a defector, should it lower, increase, or have no effect on my standing? Should my act of kindness be viewed differently if I myself was in bad standing? To tackle this question Ohtsuki and Iwasa (2004) examined a vast range of assignment systems and found eight leading strategies. All eight of these strategies share the property that cooperating with good folks is good and defecting on good folks is bad. Moreover, in terms of all leading eight strategies it is good to defect on bad folks. Empirical studies on how humans actually adjust their impressions of others can help narrow the search further. A third issue is that indirect reciprocity models assume that information quality is sufficiently high because mechanisms like gossip allow reputational information to be broadcast. But why are people motivated to tattle tale when they personally do not stand to gain from it? And why do people not manipulate information for their own gains enough to undermine the quality of reputational information and thereby unravel cooperation? A better grasp of these issues will provide a more accurate picture of how indirect reciprocity evolved, and will get us closer to describing the structure and nuances of our moral psychology.

Punishment

Cooperation can be maintained if individuals cooperate, and at some cost to themselves, punish those who defect. Such altruistic punishment can be evolutionarily stable against defectors as long as punishment imposes a larger cost on defectors than the benefit defectors acquire by free riding (Boyd et al., 2003). An advantage of punishment is that it can maintain cooperation in much larger groups than indirect reciprocity because people do not need to know the previous behavior of everyone else in the group. For instance, indirect reciprocity alone is unlikely to maintain cooperation on Turkana raids, as raiding parties comprise a few hundred men who are drawn from different segments of society and do not know each other well. This may be why the Turkana occasionally also dispense corporal punishment to cowards and deserters in raids. A problem with punishment is that it can be undermined by second-order free riders. In indirect reciprocity, the discriminator strategy benefits from refusing to help someone in bad standing, and so are evolutionarily stable against unconditional cooperators (Panchanathan and Boyd, 2004). In contrast, because punishment is costly, unconditional cooperators or second-order free riders fare better than punishers by gaining the benefit of punishment without paying the cost (Yamagishi, 1986).

Empirically, humans are motivated to cooperate and to punish those who do not, even at a cost to themselves (Gintis et al., 2005). In a classic experiment Fehr and Gachter (2002) compared players' contribution in multiple rounds of a public goods game when players could either punish or not punish shirkers in their group. When punishment was allowed, players used it even though it cost them – 84% of players punished at least once – and consequently contribution levels increased each round. In contrast, contributions decreased each round in a typical public goods experiment without punishment because conditional cooperators who start out contributing stop doing so when they notice free riders in their group. Remarkably, experiments also reveal that the motivation to engage in costly punishment extends to uninvolved third-party observers. In the third-party punishment game, an allocator receives an endowment and can allocate any amount of it to the recipient. After the allocation, a third person – the observer – can reduce her own endowment to subtract money from the allocator. Henrich et al. conducted this experiment across 15 small-scale societies (Henrich et al., 2006). In every society, there were observers who would punish unfair allocators, even though such punishment is costly to the observer and provides the observer no personal benefit.

There is evidence that punishment supports cooperation outside of laboratory experiments too. Since their inception, state societies have utilized punishment to curb norm violators and criminals, and enforce rule of law. However, given that punishment is costly to implement, it is important to know if punishment occurs in politically uncentralized societies where there are no police and courts responsible for implementing punishment. Here the Turkana provide a useful case study. Turkana warriors report that they do indeed punish cowards and deserters in raiding parties (Mathew and Boyd, 2011). A group of age-mates find the unsuspecting violator, restrain him by holding or tying him, and beat or whip him. All the while, they reprimand him for what he did, teach him how he should have behaved, and urge him to not repeat the offense in future raids. After this, the violator may be required to kill an animal from his herd for the others to eat. The Turkana view such punishment as an intervention to correct men who have erred, rather than as an opportunity to injure a person or dispense violence (Mathew and Boyd, 2014). Age-mates and other community members discuss the matter beforehand during which differing opinions may be heard about whether punishment is warranted. Some violators are only criticized verbally and warned that if they repeat it, they will face more serious sanctions.

Punishment is generally accepted to play some role in large-scale cooperation, but there are ongoing debates about how important its role is. One point of contention is how punishment compares to rewards and other forms of positive reinforcement. Punishment may not be efficient from the group's perspective (Ohtsuki et al., 2009). Because costly punishment lowers the payoff of both the punisher and punished, in some experiments it yields lower average payoffs for the group (Dreber et al., 2008). When both the option for rewarding cooperators and punishing defectors is allowed, rewarding increased contribution levels more than punishment in the repeated public goods game (Rand et al., 2009). Other experiments contradict these findings. When the time horizon is longer than in typical laboratory experiments, punishment increased average group payoffs (Gächter et al., 2008). In the study by Ule et al. (2009) where individuals had the opportunity to either help, not help, or punish a partner, players used indirect sanctions much more often than direct punishment. However, the actions of the 10% of players who were other-regarding costly punishers were necessary to make defectors have a lower payoff than other-regarding discriminators.

Another issue is whether there is sufficient empirical evidence for costly punishment outside of laboratory experiments (Guala, 2012). Some scholars contend that in the real world punishment is typically implemented through low-cost means. For instance, the Turkana first criticize cowards and deserters, and utilize corporal sanctions when such verbal sanctions do not suffice. Moreover, a group of age-mates convene to punish a violator. Such coordinated punishment is much cheaper to implement and can evolve under a less restrictive range of conditions (Boyd et al., 2010). While a solo punisher may face physical harm and retaliation, age-mates who deliver punishment in a group escape these costs. Nonetheless, in addition to the chance of potentially losing a valuable social tie, organizing corporal punishment takes time and effort. Age mates discuss the violation, gather local age-mates, plan the punishment, find the unsuspecting violator, and discipline him. These costs, albeit lower than what is faced by a solo punisher, matter when considering the evolution of a trait. So, the question of why people are motivated to bear costs to punish free riders remains a crucial open problem.

Although some theoretical solutions to the second-order free rider problem have been proposed (e.g., Boyd et al., 2003; Henrich and Boyd, 2001), more work is needed to clarify the issue. There are virtually no empirical studies of second-order free riding in small-scale societies, which raises many questions. Do societies have cultural norms against second-order free riding? Are people motivated to sanction second-order free riders, or are we motivated to reward punishers? Is third- or fourth-order free-riding a concept in any society? Thus it remains to be seen how humans solved the critical problem of second-order free riding and evolved the capacity for punishment.

Cultural Group Selection

Indirect reciprocity and punishment have explained important aspects of human cooperation and moral motivations, but they are not complete explanations. Sanctions have the unfortunate property that they can sustain a vast range of outcomes (Boyd and Richerson, 1992). Sanctions can maintain both cooperative and noncooperative outcomes. For instance mutual enforcement can maintain the norm 'You may steal your neighbor's unguarded cows to feed your family' or 'You may not steal your neighbor's unguarded cows to feed your family'. Sanctions can also maintain norms on a wide range of scales. For example, 'Do not steal your clansmen's cattle, but the cattle of the rest are for brave men to steal' or 'Do not steal your tribesmen's cattle, but the cattle of other tribes are for brave men to steal' are both group-beneficial norms, but one benefits clans, the other benefits tribes. In the domain of property rights, cooperation occurs on the scale of clans in the former society, and on the much larger scale of tribes in the latter society.

Experimental findings too confirm what casual observations tell: human societies vary widely in the norms they enforce, and these norms support cooperation at vastly different scales. Herrmann et al. (Herrmann et al., 2008) conducted a public goods game with punishment in 16 countries, and found that participants of some countries punished individuals who contributed to the public goods, while in other countries people punished noncontributors. Antisocial punishment was more common in countries with weak rule of law and low civic cooperation. Another study comparing behavior across 15 small-scale societies (Henrich et al., 2010) showed that people in societies with higher market integration and large community size are more generous in the dictator game and more willing to punish unfair offers in the ultimatum game. This suggests that large, market-based societies have norms of fairness in transient interactions that support cooperation with anonymous strangers, whereas small-scale societies lack these norms.

A complete account of human cooperation would specify the processes that select among the many possible norms that sanctions can sustain to yield the patterns of cooperation we observe across societies.

Cultural group selection (Boyd and Richerson, 1985) can explain why sanctions and norms support cooperation on the scale of cultural groups – symbolically marked groups of people with similar norms and beliefs. Social learning can maintain between-group variation in behavior despite substantial migration, making cultural group selection much more plausible than genetic group selection (Henrich and Boyd, 1998; Boyd and Richerson, 2010). Bell et al. (2009) tested this claim by estimating between-group variation in cultural and genetic traits among contemporary societies. The analysis shows that cultural FST values are orders of magnitude larger than genetic FST values, supporting the view that cultural group selection is far more plausible than genetic group selection. How do social learning mechanisms maintain sufficient between-group behavioral variation? Conformism – imitating the behavior of the majority – is one type of social learning mechanism that can evolve in a wide range of conditions that favor social learning (Boyd and Richerson, 1985; Henrich and Boyd, 1998). If people have conformist learning biases, then migrants into a group will adopt the behavioral variant of the new group, because they adopt the common variant in the new environment. A second way by which social learning biases can maintain between-group variation is by punishment (Richerson et al., 2003). If norm violations are punished, then new immigrants into a group will quickly learn to adopt the customs of the new social environment. As long as there is sufficient between-group variation, there are ample means by which cultural groups can compete: warfare, differential population growth, and greater success at being imitated by other groups. Consequently norms that benefit the cultural group can evolve.

There is some evidence that the scale of human cooperation is consistent with cultural group selection. Consider the following study among the Turkana (Mathew and Boyd, 2011). The Turkana are subdivided into two-dozen territorial sections each comprising 20 000–50 000 people. Participants from one territorial section – the Kwatela – were narrated one of two scenarios. In one scenario, two Kwatela warriors go to the neighboring Turkana territorial section – the Lukumong. They steal cows, and drive them back to Kwatela land. In the other scenario, two Kwatela warriors go to the neighboring tribe – the Toposa – and steal cows, which they drive back to Kwatela land. Participants, who were all Kwatela, were then asked a series of questions to elicit how they judged the warriors. Strikingly, although in both scenarios the Kwatela benefit, participants reacted extremely differently. They felt that the warriors who raided Lukumong cattle did something very wrong, deserved sanctions, and were undeserving of their help. In contrast, the warriors who raided Toposa cattle were praiseworthy and deserving of help. This means that Turkana norms governing warfare create benefits exactly at the scale of the cultural group – not at the smaller scale of territorial sections, and not at the larger scale that includes other tribes. The pattern of male mortality in Turkana suggests that the norms really matter. About half of adult male mortality is due to interethnic warfare, and only 1% is due to internal interpersonal violence. Thus, despite the ubiquity of firearms within Turkana settlements, the Turkana successfully suppress interpersonal violent conflict, and most of the conflict is between the Turkana and other tribes. If this pattern is general, there is a compelling case to be made that the ultrasocial unit in human societies is the cultural group. As cultural groups get bigger, so then will the scale at which humans cooperate.

Conclusion

For a long time, evolutionary theory did not explain human behavior as adequately as it explained the behavior of other animals. But recent models of cooperation have begun to integrate distinctive features of humans such as culture and moral sentiments into the evolutionary framework. Direct reciprocity, and the range of reciprocal strategies that deal with errors, have provided a useful framework for understanding pairwise relations in humans. Indirect reciprocity and punishment models have illuminated how humans evolved to cooperate in large groups of unrelated individuals. Finally, cultural group selection explains the scale of human cooperation, and why it is variable. There remain many unresolved questions: What reciprocal strategy underpins friendships? What is goodness? Why is gossip honest enough? Why are people motivated to punish? More theoretical work and empirical studies across a wide range of societies can help fill these gaps, and provide us a more detailed picture of how we became ultrasocial.

See also: Altruism and Prosocial Behavior, Sociology of; Behavioral Economics; Charity and Philanthropy, Economics of; Collective Behavior, Sociology of; Common Law; Conflict and War: Anthropological Aspects; Conventions and Norms: Philosophical Aspects; Cooperative Game Theory; Cross-Cultural Research Methods in Sociology; Cultural Evolution: Overview; Cultural Evolution: Theory and Models; Decision and Choice: Bounded Rationality; Deterrence Theory: Crime; Deterrence: Sanction Perceptions; Deterrent Effect of Police and Prisons; Economic Anthropology; Egalitarianism; Emergent Properties; Empathy: A Social Neuroscience Perspective; Equity in Law; Ethnic Conflicts; Ethnocentrism; Evolutionary Approaches to Human Behavior: Philosophical Aspects; Evolutionary Economics; Experimental Economics; Fertility Theory: Embodied-Capital Theory of Life History Evolution; Fertility Theory: Theory of Life History Evolution; Field Experiments; Field Observational Research in Anthropology and Sociology; Friendship, Anthropology of; Game Theory: Noncooperative Games; Game Theory; Gene–Culture Coevolution; Human Behavioral Ecology; Human Cognition, Evolution of; Hunting and Gathering Societies: Anthropology; Institutions; Intentionality in Language and Communication, Emergence of; Justice, Conflicts, and the Justice of Conflict Resolution; Kin Selection; Kinship, Evolution of; Kinship, Formal Models of; Law and Economics; Law and Morality: An Analytical Perspective; Law: Anthropological Aspects; Market and Nonmarket Allocation; Mind, Theories of; Morality 'East' and 'West': Cultural Concerns; Morality: Evolution of; Natural Law; Neural Foundation of Morality; Norms; Philosophy of Law; Property Rights, History of; Public Goods: International; Punishment: Comparative Politics of; Punishment: Social and Legal Aspects; Punishment; Quantitative Cross-national Research Methods; Religion:

Morality and Social Control; Rule of Law (and Rechtsstaat); Sanctions in Political Science; Social Networks; Sociality In Anthropology; Sustainability and Sustainability Science; Tradition: Social; Trust: Philosophical Aspects; Welfare: Philosophical Aspects.

Bibliography

Alexander, R.D., 1987. The Biology of Moral Systems. Aldine de Gruyter.

Allen-Arave, W., Gurven, M., Hill, K., 2008. Reciprocal altruism, rather than kin selection, maintains nepotistic food transfers on an Ache reservation. Evolution and Human Behavior 29, 305–318.

Axelrod, R., Hamilton, W.D., 1981. The evolution of cooperation. Science 211 (4489), 1390–1396.

Barkow, J.H., Tooby, J., Cosmides, L. (Eds.), 1992. The Adapted Mind: Evolutionary Psychology and the Generation of Culture. Oxford University Press, Oxford.

Bateson, M., Nettle, D., Roberts, G., 2006. Cues of being watched enhance cooperation in a real-world setting. Biology Letters 2 (3), 412–414.

Bell, A.V., Richerson, P.J., McElreath, R., 2009. Culture rather than genes provides greater scope for the evolution of large-scale human prosociality. Proceedings of the National Academy of Sciences 106 (42), 17671–17674.

Bowles, S., Gintis, H., 2011. A Cooperative Species: Human Reciprocity and Its Evolution. Princeton University Press, Princeton, NJ.

Boyd, R., 1989. Mistakes allow evolutionary stability in the repeated prisoners-dilemma game. Journal of Theoretical Biology 136 (1), 47–56.

Boyd, R., et al., 2003. The evolution of altruistic punishment. Proceedings of the National Academy of Sciences 100 (6), 3531–3535.

Boyd, R., Gintis, H., Bowles, S., 2010. Coordinated punishment of defectors sustains cooperation and can proliferate when rare. Science 328, 617–620.

Boyd, R., Richerson, P., 2010. Transmission coupling mechanisms: cultural group selection. Philosophical Transactions of the Royal Society (B) 365 (1559), 3787–3795.

Boyd, R., Richerson, P.J., 1985. Culture and the Evolutionary Process. University of Chicago Press, Chicago, IL.

Boyd, R., Richerson, P.J., 1992. Punishment allows the evolution of cooperation (or anything else) in sizable groups. Ethology and Sociobiology 13 (3), 171–195.

Camerer, C.F., 2003. Behavioral Game Theory: Experiments on Strategic Interaction. Princeton University Press, Princeton, NJ.

Clutton-Brock, T., 2009. Cooperation between non-kin in animal societies. Nature 462, 51–57.

Cosmides, L., Barrett, H.C., Tooby, J., 2010. Adaptive specializations, social exchange, and the evolution of human intelligence. Proceedings of the National Academy of Sciences 107, 9007–9014.

Dreber, A., et al., 2008. Winners don't punish. Nature 452, 348–351.

Ein-Dor, T., et al., 2010. The Attachment Paradox: How can so many of us (the insecure ones) have no adaptive advantages? Perspectives on Psychological Science 5 (2), 123–141.

Falk, A., Fehr, E., Fischbacher, U., 2008. Testing theories of fairness – Intentions matter. Games and Economic Behavior 62, 287–303.

Fehr, E., Gachter, S., 2002. Altruistic punishment in humans. Nature 415, 137–140.

Fudenberg, D., Maskin, E., 1990. Evolution and cooperation in noisy repeated games. The American Economic Review 80 (2), 274–279.

Gächter, S., Renner, E., Sefton, M., 2008. The long-run benefits of punishment science. Science 322 (5907), 1510.

Gintis, H., et al. (Eds.), 2005. Moral Sentiments and Material Interests: The Foundations of Cooperation in Economic Life. MIT Press, Cambridge, MA.

Guala, F., 2012. Reciprocity: Weak or Strong? What punishment experiments do (and do not) demonstrate. Behavioral and Brain Sciences 35 (1), 1–15.

Gurven, M., 2004. Reciprocal altruism and food sharing decisions among Hiwi and Ache hunter–gatherers. Behavioral Ecology and Sociobiology 56, 366–380.

Gurven, M., et al., 2001. Reservation food sharing among the Ache of Paraguay. Human Nature 12 (4), 273–297.

Haley, H., Fessler, D., 2005. Nobody's watching? Subtle cues affect generosity in an anonymous economic game. Evolution and Human Behavior 26, 245–256.

Hamilton, W.D., 1964. The genetical evolution of social behaviour. Journal of Theoretical Biology 7 (1), 17–52.

Henrich, J., et al., 2005. "Economic man" in cross-cultural perspective: behavioral experiments in 15 small-scale societies. Behavioral and Brain Sciences 28, 795–815.

Henrich, J., et al., 2006. Costly punishment across human societies. Science 312 (5781), 1767–1770.

Henrich, J., et al. (Eds.), 2004. Foundations of Human Sociality: Economic Experiments and Ethnographic Evidence from Fifteen Small-scale Societies. Oxford University Press, Oxford.

Henrich, J., et al., 2010. Markets, religion, community size, and the evolution of fairness and punishment. Science 327, 1480–1484.

Henrich, J., Boyd, R., 1998. The evolution of conformist transmission and the emergence of between-group differences. Evolution and Human Behavior 19, 215–241.

Henrich, J., Boyd, R., 2001. Why people punish defectors: weak conformist transmission can stabilize costly enforcement of norms in cooperative dilemmas. Journal of Theoretical Biology 208, 79–89.

Henrich, J., Henrich, N., 2007. Why Humans Cooperate: A Cultural and Evolutionary Explanation. Oxford University Press, Oxford.

Herrmann, B., Thoni, C., Gachter, S., 2008. Antisocial punishment across societies. Science 319, 1362–1367.

Hill, K., et al., 2011. Co-residence patterns in hunter-gatherer societies show unique human social structure. Science 331 (6022), 1286–1289.

Hill, K., Hurtado, M., 2009. Cooperative breeding in South American hunter-gatherers. Proceedings of the Royal Society of London B 276, 3863–3870.

Hrdy, S., 2009. Mothers and Others: The Evolutionary Origins of Mutual Understanding. Harvard University Press, Cambridge, MA.

Kaplan, H., et al., 2000. A theory of human life history evolution: diet, intelligence, and longevity. Evolutionary Anthropology 9 (4), 156–185.

Kurzban, R., 2001. The social psychophysics of cooperation: nonverbal communication in a public-goods game. Journal of Nonverbal Behavior 25, 241–259.

Leimar, O., Hammerstein, P., 2001. The evolution of cooperation through indirect reciprocity. Proceedings of the Royal Society of London B 268, 745–753.

Maner, J.K., Kenrick, D.T., 2010. When Adaptations Go Awry: Functional and Dysfunctional aspects of social Anxiety. Social Issues and Policy Review 4 (1), 111–142.

Mathew, S., Boyd, R., 2011. Punishment sustains large-scale cooperation in prestate warfare. Proceedings of the National Academy of Sciences 108 (28), 1091–6490.

Mathew, S., Boyd, R., 2014. The cost of cowardice: punitive sentiments towards free riders in Turkana raids. Evolution and Human Behavior 35 (1), 58–64.

McElreath, R., Boyd, R., 2008. Mathematical Models of Social Evolution: A Guide for the Perplexed. University of Chicago Press, Chicago, IL.

Milinski, M., Semmann, D., Krambeck, H.J., 2002. Reputation helps solve the 'tragedy of the commons'. Nature 415, 424–426.

Nowak, M., 2006. Five rules for the evolution of cooperation. Science 314, 1560–1563.

Nowak, M., Sigmund, K., 1992. Tit for tat in heterogeneous populations. Nature 355 (6357), 250–253.

Nowak, M.A., Sigmund, K., 1993. A strategy of win-stay, lose-shift that outperforms tit-for-tat in Prisoner's Dilemma. Nature 364 (6432), 56–58.

Nowak, M.A., Sigmund, K., 1998. Evolution of indirect reciprocity by image scoring. Nature 393, 673–676.

Ohtsuki, H., Iwasa, Y., 2004. How should we define goodness?—reputation dynamics in indirect reciprocity. Journal of Theoretical Biology 231 (1), 107–120.

Ohtsuki, H., Iwasa, Y., Nowak, M.A., 2009. Indirect reciprocity provides only a narrow margin of efficiency for costly punishment. Nature 457 (7225), 79–82.

Panchanathan, K., Boyd, R., 2003. A tale of two defectors: the importance of standing for evolution of indirect reciprocity. Journal of Theoretical Biology 224 (1), 115–126.

Panchanathan, K., Boyd, R., 2004. Indirect reciprocity can stabilize cooperation without the second-order free rider problem. Nature 432 (7016), 499–502.

Rand, D.G., et al., 2009. Positive interactions promote public cooperation. Science 325, 1272–1275.

Richerson, P.J., Boyd, R., 2005. Not by Genes Alone: How Culture Transformed Human Evolution. The University of Chicago Press, Chicago, IL.

Richerson, P.J., Boyd, R., Henrich, J., 2003. Cultural evolution of human cooperation. In: Hammerstein, P. (Ed.), Genetic and Cultural Evolution of Cooperation. MIT Press, Cambridge, MA.

Smith, J., Szathmary, E., 2000. The Origins of Life: From the Birth of Life to the Origin of Language. Oxford University Press, Oxford.

Sugiyama, L.S., 2004. Illness, injury, and disability among Shiwiar forager-horticulturalists: implications of health-risk buffering for the evolution of human life history. American Journal of Physical Anthropology 123 (4), 371–389.

Trivers, R.L., 1971. The evolution of reciprocal altruism. The Quarterly Review of Biology 46 (1), 35–57.

Ule, A., et al., 2009. Indirect punishment and generosity toward strangers. Science 326 (5960), 1701–1704.

van Veelen, M.J., et al., 2012. Direct reciprocity in structured populations. Proceedings of the National Academy of Sciences 109 (25), 9929–9934.

Xue, M., Silk, J., 2012. The role of tracking and tolerance in relationship among friends. Evolution and Human Behavior 33 (1), 17–25.

Yamagishi, T., 1986. The provisioning of sanctioning as a public good. Journal of Personality and Social Psychology 51, 100–116.