

the increase in species richness with island area is much more gradual than in a comparable continental region (that is, many island species–area relationships have much smaller exponents than in continental regions).

Distribution–abundance relationships and related patterns are a consequence of asymmetries in ecological attributes among species such that some species are able to find more resources across geographic space than others. Thus, within metacommunities, the relative abundances of species will also be asymmetric, and such asymmetries in relative abundances in metacommunities are preserved even under short-term ecological drift. When a collection of local communities are sampled using sampling units of the same size (thus controlling for area effects), the asymmetries in relative abundances of species within the metacommunity will ensure that some species show up in samples more often than others (Fig. 2). Abundances of species at local scales will also reflect these asymmetries; hence, species that show up more often in samples will also tend to be represented by more individuals within those samples.

When examining species diversity along environmental gradients, the sampled geographic region may represent only a subset of sites contained within the metacommunity of the region. Since there is spatial autocorrelation in environmental conditions, sites that have more resources will typically have more species, and these sites will tend to be clustered together along the gradient. In more extensive samples, such as samples of species diversity

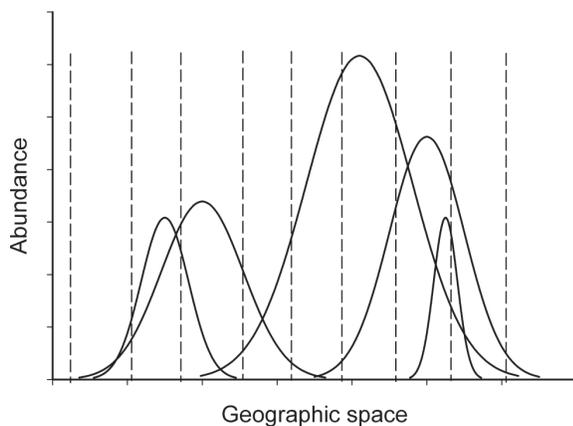


FIGURE 2 A schematic diagram showing how asymmetries in distributions of different species across geographic space result in differences in species diversity among a collection of sampling locations. Each unimodal curve represents the geographic distribution of a single species. Note that some species are found across a greater span of geographic space than others. This would be a consequence of the fact that some species will have attributes that will allow them to use a wider range of environmental conditions than others. The dashed lines represent geographic locations where a survey is taken. Surveys are assumed to be of standard size.

with latitude, many metacommunities may be sampled, and the spatial autocorrelation among resources may be very low for regions at different extremes along the gradient. In such cases, the trends in species diversity observed across space must reflect differences in the capacity of the environment to support species.

SEE ALSO THE FOLLOWING ARTICLES

Allometry and Growth / Diversity Measures / Metacommunities / Neutral Community Ecology / Species Ranges

FURTHER READING

- Blackburn, T.M., and K.J. Gaston, eds. 2003. *Macroecology: concepts and consequences*. Proceedings of the 43rd Annual Symposium of the British Ecological Society (17–19 April 2002). Malden, MA: Blackwell Publishing.
- Brown, J. H. 1995. *Macroecology*. Chicago: University of Chicago Press.
- Gaston, K. J., and T.M. Blackburn. 2000. *Pattern and process in macroecology*. Malden, MA: Blackwell Science.
- Hubbell, S. P. 2001. *The unified neutral theory of biodiversity and biogeography*. Princeton, NJ: Princeton University Press.
- Huston, M. A. 1994. *Biological diversity*. Cambridge, UK: Cambridge University Press.
- Maurer, B. A. 1999. *Untangling ecological complexity: the macroscopic perspective*. Chicago: University of Chicago Press.
- Maurer, B. A. 2009. Spatial patterns of species diversity in terrestrial environments. In S.A. Levin, ed. *Princeton guide to ecology*. Princeton, NJ: Princeton University Press.
- Rosenzweig, M.L. 1995. *Species diversity in space and time*. Cambridge, UK: Cambridge University Press.

CONTROL THEORY

SEE OPTIMAL CONTROL THEORY

COOPERATION, EVOLUTION OF

MATTHEW R. ZIMMERMAN, RICHARD MCELREATH, AND PETER J. RICHERSON

University of California, Davis

Cooperation occurs when individuals act together for beneficial results. Most evolutionary theory of cooperation addresses the evolution of altruism, the most difficult type of cooperation to explain. However, observed cooperation can also result from environments favoring other kinds of cooperation, including mutualism and coordination. Theoreticians model the evolution of cooperation by accounting for the effects of cooperative behavior at different levels of analysis, including the genes, the individual,

and the group. Different mechanisms have been proposed as ones that encourage cooperative behavior, including limited dispersal, signaling, reciprocity, and biased transmission. All of these mechanisms are based on cooperators assortatively interacting with other cooperators.

TYPES OF COOPERATION

Mutualism

Mutualistic behavior is helping behavior for which the producer's costs are smaller than the producer's benefits. Intraspecific mutualisms lack clear incentives for individuals to withhold cooperation, and so natural selection always favors cooperation over noncooperation. (Community ecologists reserve the term mutualism for interspecies relationships, though here, as in evolutionary game theory generally, we use the term primarily for intraspecific interactions.) To see how mutualisms evolve, imagine a population of wolves that can either pair up to hunt deer cooperatively or hunt deer noncooperatively. (Of course, wolves can cooperatively hunt in larger groups, but the underlying logic of mutualisms holds even if we restrict our analysis to two individuals.) If either of the wolves hunts individually, they cannot bring down large game and go hungry. However, if the wolves hunt together, they can capture a deer and eat three units of meat.

How would cooperation evolve in a population of wolves made up of cooperative hunters and noncooperative hunters? We will assume that wolves that are more successful hunters are more likely to produce offspring. The potential payoffs to the wolves in this situation are shown in Figure 1A. Both wolves have a choice of cooperative hunting (C) and noncooperative hunting (N). The numbers in the grid show the payoffs for each combination of behaviors, with the payoff to the Wolf A first, followed by the payoff to Wolf B. Notice that a wolf only has a chance at getting meat if it tries to hunt cooperatively. Thus, wolves that hunt cooperatively will tend to produce more offspring and the trait for cooperative hunting will tend to spread in the population.

The line in Figure 1B shows the adaptive dynamics of this wolf population. The arrows represent the direction of natural selection and show that, no matter where the mixture of hunting strategies starts, selection will tend to push the population toward a stable cooperative equilibrium, represented by the solid circle. The open circle at the other end of the line represents a noncooperative equilibrium where there are no cooperative wolves and they cannot reproduce. However, this equilibrium is

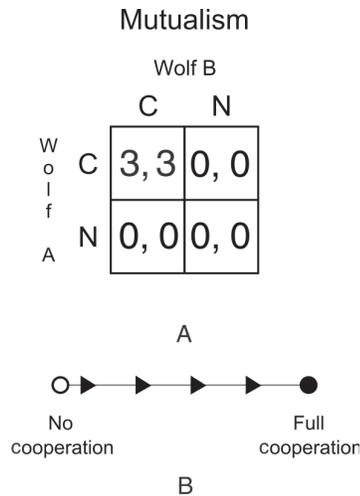


FIGURE 1 (A) A simple game illustrating mutualistic hunting in a hypothetical wolf population. Wolves get meat only if they hunt cooperatively, thus the benefits of cooperation are always positive. (B) The adaptive dynamics of the wolf population when individuals are randomly paired. Because it is always better to cooperate, natural selection increases the frequency of cooperative hunting away from the unstable noncooperative equilibrium to the stable equilibrium at full cooperation.

unstable since any introduction of cooperative wolves, through genetic mutation or migration from other populations, would allow natural selection to push the population to the stable equilibrium.

In some models of mutualisms, since cooperative individuals always achieve greater benefits than noncooperative individuals, the evolution of cooperative behavior does not present much of a puzzle and does not require special explanations for why they persist in a population. However, the origins of mutualisms can be more difficult to explain, as initial benefits may not be sufficient to exceed individual costs. For example, if in a proto-wolf population hunters are mainly solitary, then it stands to reason that the skills necessary for successful cooperative hunting will not exist. Therefore, early cooperative hunters are not assured the large gains that may arise later, once cooperation gets a foothold through other mechanisms. In this way, other types of cooperation, such as coordination (discussed next), can evolve over time into mutualisms, once cooperation is stabilized through other means.

Coordination

In coordination contexts, natural selection favors common behavior, whether it is cooperative or noncooperative. For example, imagine that our population of wolves lives in an environment with both deer and rabbits. Wolves now have two strategies: hunt deer cooperatively, or solitary hunt rabbits. If a wolf and its partner hunt cooperatively,

they manage to catch a deer with three units of meat apiece as before. If, however, one of the wolves hunts deer and the other hunts rabbits, the deer hunter fails to bring down the game, but the rabbit hunter catches two rabbits for two units of meat. However, if both wolves hunt rabbits individually, they are competing for the easy game and only receive one unit of meat each. The payoffs for this situation are shown in Figure 2A.

In this wolf population, the direction of natural selection depends on the current level of cooperative hunting. For example, if all the other wolves in the population are cooperatively hunting deer, a rogue rabbit hunter will have less meat and will therefore be less reproductively successful than the rest of the population. Thus, natural selection works against rabbit hunters. However, if all the other wolves in the population are hunting rabbits, then a lone deer hunter will not find a cooperative hunting partner and starve. Natural selection, in this case, works against the deer hunters.

The line in Figure 2B illustrates the adaptive dynamics of this wolf population. There are two stable equilibria. If deer hunters are more common in the population, natural selection drives the population toward the cooperative deer-hunting equilibrium. If rabbit hunters are more common, natural selection drives the population toward

the noncooperative rabbit-hunting equilibrium. These equilibria are stable because the introduction of a non-conforming wolf to the population will have lower payoffs than the rest of the population obtains and thus be outcompeted. Because, in coordination, the equilibrium preferred by natural selection depends on the frequency of behaviors in the population, it is called frequency dependent. (There is also an unstable equilibrium where there are exactly as many deer hunters as rabbit hunters, but this equilibrium will not persist, as natural selection will tend to drive the population away from the center as soon as chance factors cause this perfect balance to be broken.)

This situation shows why coordination can be a more complicated type of cooperation than mutualism. A population of rabbit-hunting wolves would be better off if they just cooperated and hunted deer, in that they would all obtain three units of meat instead of one, but natural selection discourages the population from moving toward the cooperative equilibrium. In resolving this dilemma of coordination, evolutionary theorists focus on equilibrium selection, how a population might move, through the processes of natural selection or drift, from the noncooperative to the cooperative equilibrium. Once a population moves to the cooperative equilibrium, it is less difficult to explain its maintenance. However, our hypothetical wolf population faces a particularly hard case of equilibrium selection since trying to hunt deer without a partner has zero payoff when most wolves are hunting rabbits. It is thus much riskier for a wolf to switch from a noncooperative to a cooperative behavior than the other way around, and the population might spend more time in the noncooperative equilibrium. A population transitioning from the noncooperative equilibrium to the cooperative equilibrium may require mechanisms similar to the evolution of altruism, as described below.

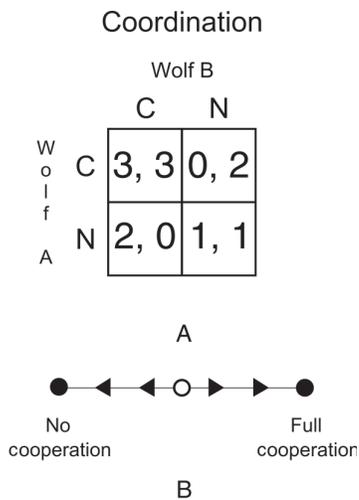


FIGURE 2 (A) A simple game illustrating a coordination scenario in a hypothetical wolf population. When individuals are randomly paired, wolves can expect to get the most meat when they have the most common behavior in the population. (B) The adaptive dynamics of the randomly paired wolf population. When cooperation is common, natural selection drives the frequency of cooperative behavior toward the stable equilibrium at full cooperation. However, when cooperation is rare, natural selection drives the frequency of cooperative behavior toward the stable equilibrium where none of the wolves cooperate. The location of the unstable equilibrium in this model depends on the payoffs in (A) and the assumption that the fitness effect of meat consumption is linear.

Altruism

A helping behavior is altruistic if the individual fitness costs of the behavior exceed the individual fitness benefits. Altruism is a context in which natural selection on individuals appears to favor noncooperation, even though average fitness across the population is assumed to increase as more individuals cooperate. Altruism is thus a strategic setting that directly opposes the interests of individuals and the interests of groups. Again imagine our hypothetical population of wolves, now living in an environment where rabbits are larger and more abundant. Now, if a wolf hunts rabbit and its partner hunts deer, it gets four units of meat while its partner gets zero. If they both hunt rabbits, they each get two units

of meat. If they both hunt deer cooperatively, they each get three units of meat. These payoffs are reflected in Figure 3A. As might be expected, if the benefits of hunting rabbits are higher, then natural selection more strongly favors hunting rabbits over cooperatively hunting deer. In this case, hunting rabbits is always a better strategy for an individual wolf than hunting deer. If a wolf's partner hunts rabbit, also hunting rabbits nets it two units of meat compared to going hungry when hunting deer. If a wolf's partner hunts deer, rabbit hunting nets it four units of meat compared to three for hunting deer. Since hunting rabbits always produces higher payoffs for the individual, natural selection should push the population toward a stable equilibrium of noncooperation, as shown in the line under the payoff matrix in Figure 3B.

Notice that if all wolves in the population hunt rabbits, they all gain two units of meat. But if they all cooperatively hunted deer, they would all have three units. Thus, the equilibrium favored by natural selection is not where the wolf population has the highest total payoff. This type of scenario is the popular model of altruism called a "Prisoner's Dilemma" (named after a hypothetical dilemma faced by prisoners who have to decide whether to rat out their accomplices during police questioning). In the Prisoner's Dilemma (PD), the cooperative strategy is an example of an altruistic act because it requires giving up the individual benefits of hunting rabbits in order to try to achieve the collective benefits of hunting deer.

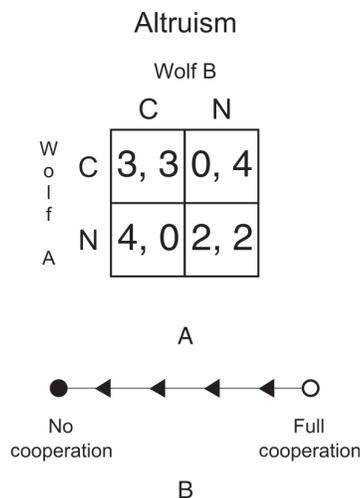


FIGURE 3 (A) A simple game illustrating the difficulty in explaining the evolution of altruistic behavior. Wolves always have a higher payoff when they behave noncooperatively, regardless of what their partner does. (B) The adaptive dynamics in a randomly paired wolf population. Because the individual wolf always gets more meat when it behaves noncooperatively, natural selection drives the frequency of cooperative behavior to zero.

Hunting rabbits is called "defection" because the rabbit hunter does best when it exploits a partner's efforts to hunt cooperatively and forces the cooperative hunter to starve. Starvation, in this scenario, is called the "sucker's payoff."

Because noncooperation appears to be the only stable evolutionary equilibrium in the PD, the conditions under which altruism can evolve in situations that resemble PD are more challenging than trying to understand the evolution of coordination or mutualistic behavior. Because altruism is the most difficult type of cooperation to explain, the evolution of altruism has received by far the most theoretical attention. In fact, many biologists think of altruism as synonymous with cooperation. It is thus important for field ecologists interested in testing theory not to assume that all observed cooperative behavior is altruistic, since this implies evolution has solved a harder problem than it actually has. For example, to understand the evolutionary history of cooperative hunting in our hypothetical wolf population, it would be important to determine and quantify the other hunting options available. However, in practice, quantification of the costs and benefits of animal behavior is often difficult, especially theoretically relevant but unobserved behavior.

Threshold Cooperation

Threshold cooperation is a strategic context in which cooperation is favored when rare but selected against when common. There is a threshold frequency of cooperation at which selection changes direction. In all of the above examples, our hypothetical wolf pack always evolves toward either full cooperation or full noncooperation. This occurs because each of the simple game models favors pure types. But many populations are made up of mixed populations of both cooperators and noncooperators. For a simple model of how this can happen, imagine that to successfully take a deer, our hypothetical wolves need to first find a deer herd. Although both wolves are needed to bring down a deer, they can split up to search more efficiently, and the wolf that first spies a group of deer can signal the other to join in the kill. If they are successful in the kill, the wolves each get a payoff of three units of meat as above. However, one or both of the wolves might shirk their searching duties and, instead, spend the time hunting rabbits and leaving their partner to find the herd. If both wolves shirk searching, they never find deer and each get one unit of meat from hunting rabbits. If only one wolf searches, it is less efficient and it metabolizes, on average, the equivalent of one unit of meat, receiving a net payoff of two units. The other wolf receives one unit

of meat from hunting rabbits and three from the deer, for a total of four units. The payoffs for this scenario are in the game matrix in Figure 4A.

This type of game is variously called the hawk–dove or snowdrift game in biology or the game of chicken in economics. These games share the feature of favoring cooperation when cooperation is rare but favoring noncooperation when cooperation is common. In this scenario, if all the wolves are cooperative searchers, then wolves that shirk cooperative searching acquire more meat (four units to two) and thus leave more offspring. So shirking behavior should spread. However, if none of the other wolves search, a wolf that does should acquire more meat than all of the noncooperative wolves except for its partner (two units to one). So searching behavior would spread. If noncooperative shirking spreads when it is rare, and cooperative searching spreads when it is rare, natural selection pushes the population away from homogenous equilibria. Figure 4B shows the adaptive dynamics of this behavior. The population moves to a stable equilibrium where there are both cooperators and noncooperators.

This game can be seen as a mild version of altruism. Behaving cooperatively is mutualistic when cooperation is rare in a population. However, cooperation becomes altruistic when it is already common. The challenge for the theorist is explaining how cooperation might evolve

above what is expected at the equilibrium depicted in Figure 4B. The challenge for the empiricist is to determine, given cooperation observed in a population, if it is above the level expected given cooperation’s individual payoffs.

MODELING THE EVOLUTION OF COOPERATION

When modeling the evolution of cooperation, the modeler must account for all of cooperative behavior’s effects on the fitness of possibly many individuals. There are two dominant and mathematically equivalent styles of analysis for accomplishing this objective: inclusive fitness and multilevel selection. The inclusive fitness approach accounts for fitness effects from the perspective of individual genes, focusing on the effects of genes coding for cooperative behavior on related individuals, who may also share copies of these genes. The multilevel selection approach instead distinguishes fitness effects at the individual and the group level—for example, examining how some packs of cooperative wolves might outcompete packs of noncooperative wolves.

Modeling in either of these approaches will yield equivalent answers, as long as the modeler properly accounts for all fitness effects. For example, as long as the modeler is careful, modeling the evolution of cooperation at the gene level will give the same result as modeling cooperation as a tug-of-war between individual and group levels. Thus, old debates concerning the “correct” style of analysis have largely faded in biology. A modeler instead selects styles of analysis that are most useful to understanding a particular question or that are the most mathematically convenient. For example, modeling individual-level selection could be the most useful approach if populations lack social structure and close relatives are widely dispersed. In that situation, the modeler may only need to understand a behavior’s effect on the “classical fitness” of an organism, or its direct number of offspring. When these assumptions are relaxed, analyzing fitness at other levels of analysis might be more useful.

Kin selection models adjust individual fitness to be a special sum of the negative effects of behavior on self and the positive effects on relatives. If an individual behaves altruistically, the cost of the behavior decreases its classical fitness. However, the benefits of altruism increase the fitness of others, and if these others are genetically related, they may also have copies of the altruism gene. Since kin share genes from a recent common ancestor, genes favoring altruistic behavior toward kin can increase in the next generation as long as the net

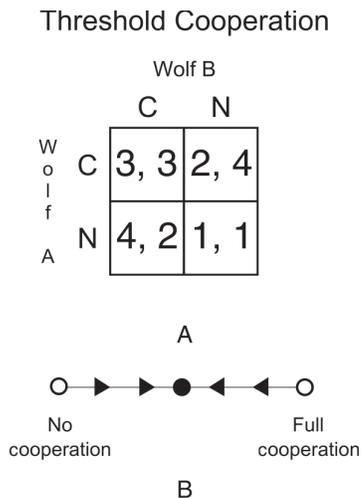


FIGURE 4 (A) A simple game illustrating threshold cooperation. Here wolves have higher payoffs when their behavior is different than the majority of other wolves in the population. (B) In this scenario, when cooperation is rare, natural selection increases its frequency in the population. When cooperation is common, natural selection decreases its frequency in the population. This creates a stable equilibrium of both cooperators and noncooperators in the wolf population. The location of the stable equilibrium in this model depends on the payoffs in (A) and the assumption that the fitness effect of meat consumption is linear.

benefits of altruism are high enough relative to the cost. For example, if a gene encourages altruistic hunting in wolves, when a wolf altruistically hunts with a sibling who shares the gene, the gene might be more likely to spread to the next generation even if altruistic behavior hurts the fitness of the original wolf. Kin selection models like this are said to consider inclusive fitness, which includes both an organism's classical fitness and the discounted fitness of its kin.

The most famous kin selection model, called Hamilton's Rule, was proposed by and named after biologist William Hamilton. It states that cooperation should spread in a population when

$$rb > c.$$

Here, c is the cost of an altruistic act to the altruistic individual, b is the benefit of the act to a target of the altruistic act, and r is the degree of relatedness between the altruist and the target of altruism. This equation implies that the less related you are to individuals you interact with, the less likely you are to behave cooperatively. Empirical work in understanding altruism in nature has been focused on trying to estimate r , b , and c for specific behaviors and determine how well they conform to Hamilton's rule. For example, a researcher could genetically sample a population of wolves to estimate r and then try to quantify the reproductive costs and benefits of cooperative hunting. With modern genetic techniques, estimating r is often much easier than estimating b or c . Unfortunately, this has led many researchers to ignore b and c when testing Hamilton's Rule, but as we have seen, determining the costs and benefits of cooperation are important for whether organisms are behaving mutualistically, altruistically, or coordinating their behavior.

In contrast, the multilevel selection approach divides the consequences of natural selection into within-group (individual) and between-group levels. For example, wolf packs with more altruistic hunters can achieve higher payoffs, increase in number, and displace wolf packs with less altruistic hunting. However, within a pack, altruistic wolves will have lower fitness than nonaltruistic wolves. Thus, selection at the group level tends to increase altruistic behavior, and selection at the individual level tends to decrease it. Any model of natural selection, whether it is about cooperation or not, can be split up in this way.

Which of these levels, individual or group, dominates depends on the magnitude of fitness effects at each level, as well as on how much variation is present at each level. For example, if all groups are identical in their levels of altruism, then selection at the group level cannot change

the frequency of altruistic genes in the population. Similarly, if groups are different in their levels of altruism but individuals within groups are all identical to others within their own groups, selection at the individual level cannot change the frequency of altruism genes.

Since group membership is often a good predictor of relatedness, kin selection models can always be reformulated as multilevel selection models, and vice versa. The choice of model depends primarily on which formulation is easier to construct, is more informative to the research question, is more open to empirical analysis, or better accounts for the relevant mechanisms that encourage cooperation in the population. Note, however, that both approaches, inclusive fitness and multilevel selection, are steady-state approximations of true evolutionary change. They are equivalent, but not exhaustive, ways of modeling the evolution of cooperation.

MECHANISMS OF POSITIVE ASSORTMENT

Much of the effort expended to understand the evolution of altruism has been to specify various mechanisms that favor its evolution or maintenance. A mechanism in this context is a strategy for achieving positive assortment of altruistic strategies. However one chooses to model the evolution of altruism, it cannot evolve unless individuals carrying altruistic genes tend to direct the benefits of altruism to other individuals carrying altruistic genes. For example, imagine that in our pack of wolves, altruistic hunters could see into the genomes of other wolves and were thereby able to pair only with other altruistic hunters. Without the risk of defection, altruistic hunters would have higher payoffs and increased numbers of offspring relative to nonaltruistic wolves interacting with other nonaltruists. While wolves cannot actually peer into one another's genomes, animals have other mechanisms that increase positive assortment.

Limited Dispersal

One mechanism that can generate positive assortment among altruists is when organisms do not disperse far from where they are born. When this occurs, altruists have greater opportunity to interact with related altruists, because kin will share space. For example, wolves may be more likely to live in packs with relatives, because many pups remain in their natal groups. This means that individuals will naturally find themselves in groups with close kin, creating positive assortment, without any need for the ability to know who their relatives are.

However, limited dispersal alone can also enhance competition among relatives and reduce altruism. Understanding

the full effects of limited dispersal also requires an understanding of how resources are distributed. If resources are locally limited, then an individual's primary reproductive competitors may also be its group mates. This localized resource competition can cancel or reverse the ability of limited dispersal to favor the evolution of altruism. It is therefore difficult to predict altruistic behavior in organisms based solely on measurements of dispersal. What is required is that altruists positively assort for social behavior but avoid one another when competing for resources.

Signaling

Signaling has also been proposed as a mechanism for generating assortment among altruists. If altruists can signal their altruism to each other, they can coordinate their interaction and avoid nonaltruists. These signals are sometimes called green beards after a hypothetical mechanism by which all altruists grow and display green beards and always help others who display green beards. Altruists can thus recognize and cooperate with each other, creating the required positive assortment. However, green beard-type signals are probably very rare in nature. It is hard for organisms to evolve green beards unless there is a way to keep nonaltruists from copying the signal. If we imagine a population of green beard altruists cooperating with each other, any individual who can display the green beard but avoid costly assistance to other green beards will have high fitness, which will eventually decouple the signal from the altruistic behavior. Because of this problem, altruistic behavior based on these signals are generally considered short-lived phenomena in complex organisms. However, in some simple organisms such as bacteria, there are possible empirical examples of single genes that both create signals and contingently help others producing the signal. Another exception might be the use of socially learned signals, such as human dialect and languages, that are so complex that outsiders cannot easily fake them.

A more commonly discussed signaling mechanism that supports altruistic behavior is kin recognition. Instead of organisms using a universal signal for altruistic behavior, organisms can recognize kin by shared traits and preferentially cooperate with them. While green beard mechanisms attempt to signal the altruism allele itself, kin recognition mechanisms attempt to recognize close kin who may share the altruism allele. If altruistic wolves remember their littermates and their littermates are genetically related to them, then the littermates that interact with one another will positively assort for

altruism. There are kin recognition strategies that use uninherited life history events correlated with kinship, such as sharing a mother, and alternative strategies that use biologically inherited markers such as facial similarity. However, current models suggest that kin recognition by inherited markers can be difficult to evolve.

Reciprocity

Reciprocity is a mechanism by which individuals can use a history of past interactions to predict an individual's probability of altruistic behavior. If wolves remember which of the members of their pack hunted altruistically in the past, they may choose to hunt with them in the future. A common framework for modeling reciprocal altruism is the iterative Prisoner's Dilemma (IPD). In the IPD, individuals play the Prisoner's Dilemma repeatedly with the same partner. Individuals are able to remember their past interactions and use that information to decide on their behavior in the next round.

A famous instance using the IPD to understand the evolution of reciprocity was two tournaments run by Robert Axelrod. Axelrod allowed participants from all over the world to submit candidate strategies for playing the IPD. The winning strategy for both tournaments, called Tit-for-Tat (TFT), was also the simplest. The strategy cooperated on the first turn and copied its opponent's behavior in the previous round on any subsequent turn. In other words, it would cooperate as long as its opponent was altruistic and defect as long as its opponent was nonaltruistic. In effect, it was a strategy that used a simple form of reciprocity based on its memory of one previous interaction.

Because of the success of TFT in Axelrod's tournament, many people think of it as the best way to play an IPD. However, later work has confirmed that TFT is not a robust strategy. TFT is not evolutionarily stable and can be reduced in frequency in a population through invasion by other cooperative strategies. For example, in a population made up only of TFT-playing individuals, a strategy that always cooperates (ALLC) will have the same payoffs as TFT, and natural selection will not act against it. In fact, if there is any cost to the cognition and memory of TFT, ALLC will have the advantage. Once ALLC becomes common enough, nonaltruistic strategies can take over the population by exploiting its lack of contingent defection. Thus, TFT is not always a successful strategy in the IPD.

TFT is not alone in its failings, however. In general, there can be no master reciprocity strategy, because many forms of contingent aid can coexist and rise and fall in

a population, depending upon the details of the model. What can be said, however, is that when animals can identify individuals and remember past interactions, some form of reciprocity that makes altruism contingent upon past behavior can often evolve. Successful strategies will often be “nice,” in the sense that they begin by cooperating in hopes that other individuals are also reciprocal.

Although the IPD model of reciprocity relied on an individual’s memory of its past interactions with others, this might not be a useful mechanism in populations where repeated interactions with the same individual are rare. In this case, individuals can obtain similar information about another’s probability of altruistic behavior from the other’s interactions with third parties. This mechanism for assortment is called indirect reciprocity. A downside to indirect reciprocity as an assortment mechanism is that for it to work as well as regular pairwise reciprocity, it requires that reputation about third parties be as accurate as direct information one obtains from personal experience. This constraint likely limits its evolution to organisms with high-fidelity communication, such as humans.

Culture and Social Learning

Thus far, the mechanisms we have discussed to explain altruistic behavior apply to any inheritance system, and for many species the most important inheritance system is genetic transmission. However, the behavior of some species, especially humans, is heavily influenced by socially learned information, or culture. The mechanisms by which cultural inheritance can produce altruistic behavior are similar to genetic inheritance. However, because culture can spread more quickly than genes to many different individuals, large groups may be much more culturally related than genetically related. This allows for altruism to evolve culturally in much larger groups than is possible with genetic transmission.

This is especially true when cultural transmission is adaptively biased. For example, under many situations it makes sense for individuals to learn the most common behaviors in their group (called a conformist bias). This allows them to quickly adapt to local conditions or contexts where coordination is important. If conformist biases are common in a species, this maintains strong cultural relatedness within a group by discrimination against learning the traits of rare migrants, which strengthens the effect of selection for altruism at the group level.

SEE ALSO THE FOLLOWING ARTICLES

Adaptive Dynamics / Behavioral Ecology /
Evolutionarily Stable Strategies / Game Theory

FURTHER READING

- Axelrod, R. 1984. *The evolution of cooperation*. New York: Basic Books.
- Clutton-Brock, T. 2009. Cooperation between non-kin in animal societies. *Nature* 462: 51–57.
- Dugatkin, L. A. 2006. *The altruism equation*. Princeton: Princeton University Press.
- Frank, S. A. 1998. *Foundations of social evolution*. Princeton: Princeton University Press.
- Lion, S., V. Jansen, and T. Day. 2011. Evolution in structured populations: beyond the kin versus group debate. *Trends in Ecology and Evolution* 26: 193–201.
- Maynard Smith, J. 1982. *Evolution and the theory of games*. Cambridge, UK: Cambridge University Press.
- Maynard Smith, J., and E. Szathmáry. 1995. *The major transitions in evolution*. Oxford: Oxford University Press.
- McElreath, R., and R. Boyd. 2007. *Mathematical models of social evolution*. Chicago: University of Chicago Press.
- Richerson, P., and R. Boyd. 2005. *Not by genes alone*. Chicago: University of Chicago Press.
- Skyrms, B. 2003. *The stag hunt and the evolution of social structure*. Cambridge, UK: Cambridge University Press.
- Sober, E., and D. S. Wilson. 1999. *Unto others*. Cambridge, MA: Harvard University Press.